Closing the Structural Gaps: Marginal Compositional Consistency and Algebraic Closure in Fact-Producing Universes

Keith Taylor *VERSF Theoretical Physics Program*

For the General Reader

Imagine trying to understand why the universe follows the rules of quantum mechanics rather than some other set of rules entirely. One approach is to ask: what rules must *any* universe follow if it is to produce definite, irreversible events at all? A universe in which nothing ever really settles — no measurement ever lands on an outcome, no coin ever stops spinning — is not a universe that contains facts. And a universe without facts is not one we could ever observe or inhabit.

In earlier work, we showed that this single requirement — that the universe must be capable of producing facts — forces a surprising amount of structure. Possibilities must be able to cancel each other out before a fact is formed. Situations that are physically equivalent must be treated the same way. And the mathematics describing unresolved possibilities must have a very specific algebraic form. Taken together, these constraints point uniquely to the framework of complex quantum mechanics.

However, two steps in that argument relied on assumptions we stated but did not fully justify. The first was about how probabilities change when you push a physical system toward a definite outcome: we assumed that what happens at one location does not depend on hidden information about other locations. The second was that every physical possibility has an exact mirror image that cancels it — not just some possibilities, but all of them without exception.

Both assumptions felt right, but feeling right is not a proof.

This paper provides the proofs. For the first, we show that local updating is not an independent postulate — it falls out automatically once you are precise about what it means to observe something in a world governed by facts. For the second, we show that if even one possibility existed that could not be cancelled, it would leave a permanent, detectable fingerprint on the physical world without ever becoming a fact itself. That would violate the basic rule that only facts carry physical meaning — so no such uncancellable possibility can exist.

The argument for the second point is the more technically involved of the two, and we are careful to identify exactly where it rests on an additional principle — one we name and state precisely rather than hide inside the proof. Whether that principle can itself be derived from something more basic is the main open question we leave for future work.

Abstract

Previous work in the VERSF reconstruction programme identified two insufficiently grounded structural assumptions: (A5) locality of composition at the level of configuration weights, and (S2) the existence of universal additive inverses in the superposition algebra. This paper derives (A5) and shows that (S2) follows under a more precisely targeted and physically motivated closure principle.

For (A5): we introduce *marginal compositional consistency* (MCC), strictly weaker than (A5), and derive it from admissibility via a named intermediate step — the *operational congruence principle* (Proposition 3.4) — which establishes that admissible operations act on observational equivalence classes. MCC follows immediately.

For (S2): we derive universal additive inverses from pointwise cancellability together with *internal admissible closure* (IAC) — a narrower replacement for the broad reversible-closure condition of earlier versions. The contribution is not a derivation of (S2) from fact-production alone; it is a reduction of the old broad closure requirement to a more sharply delimited and physically motivated condition, from which (S2) follows. A preliminary Faithfulness Lemma establishes that observational nullity implies algebraic nullity — via the admissibility quotient declared globally in Section 2 — and is used transparently throughout the main chain. IAC enters at one step only and yields a cancellation partner $\chi \in \mathcal{P}$ with $r + \chi \cong 0$; non-negativity then forces $r \cong 0$; faithfulness lifts this to $r = 0$, contradicting non-nullity; and Corollary 4.11 lifts the conclusion to additive inverses in A.

Table of Contents

---

1. Introduction

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

The derivation depended on two assumptions that were physically motivated but not derived from the core premises:

- **(A5) Locality of composition**: the update of weight $w(\psi, \lambda)$ at configuration λ under composition with a definite alternative depends only on the prior weight $w(\psi, \lambda)$ and the composition parameter t.
- **(S2) Universal additive inverses**: for every pre-factual state ψ there exists $(-\psi)$ such that $\psi + (-\psi)$ equals the null alternative.

The original paper acknowledged both as genuine additional structural assumptions and listed (A5) as the most important open question. This paper addresses both.

For (A5): the full locality condition is stronger than necessary. What the Cauchy functional equation argument requires is the strictly weaker condition *marginal compositional consistency* (MCC). The derivation proceeds through an explicit intermediate step: Proposition 3.4 (operational congruence) establishes that admissible dynamics must act on observational equivalence classes, because any operation mapping equivalent inputs to distinguishable outputs would render physically meaningful a distinction not grounded in commitment-traceable facts. MCC is then an immediate consequence. Note that MCC and (A5) coincide on the two-alternative sector where the Cauchy argument is applied (Remark 3.12), so deriving MCC rather than (A5) does not alter any downstream result. The gain is conceptual: MCC eliminates an implicit cross-configuration independence assumption and establishes that the Cauchy argument requires no such assumption on the full pre-factual sector.

For (S2): the contribution of this paper is a reduction, not a full derivation. We show that the broad reversible-closure condition (RC) of earlier versions can be replaced by *internal admissible closure* (IAC) — a precisely stated, narrower principle covering only observationally active, internally realised affine sub-contributions. (S2) is derived from IAC via a four-lemma chain. The strongest accurate claim is: we have reduced the old broad closure requirement to a condition restricted to the specific class needed by the proof, with explicit physical motivation; whether that condition itself follows from (A1)–(A2) alone is the primary open question.

What this paper does not do. Metric homogeneity (A3), commutativity (C5), and the imported distinguishability geometry results are unchanged. IAC is not derived from (A1)–(A2) alone.

---

2. Background and Notation

We work within the framework of the parent paper. The basic structure is:

- A set **Λ** of distinguishable configurations with distinguishability relation $\delta: \Lambda \times \Lambda \to [0,1]$.
- A pre-factual sector $\boldsymbol{\mathcal{P}}$ of pre-commitment states with reversible dynamics.
- A weight function **w**: $\mathcal{P} \times \Lambda \to [0,1]$ with $\sum_{\lambda} w(\psi, \lambda) = 1$ for all $\psi \in \mathcal{P}$.
- A one-parameter family of composition rules $\star_t$ ($t \in [0,1]$) on $\mathcal{P}$, satisfying conditions (i)–(vi) of Proposition 3.6b of the parent paper. The family includes the identity operation at t = 0: $\psi \star_0 \delta_{\mu} = \psi$ for all ψ, μ.
- An irreversible commitment map **φ**: $\mathcal{P} \to \Lambda$.
- The **null alternative 0** $\in \mathcal{P}$: the unique state with $w(0, \lambda) = 0$ for all $\lambda \in \Lambda$.
- The **admissibility constraint** (A2): only distinctions traceable to recordable facts carry physical meaning.
- **(A2c) Admissibility is closed under composition**: the composition of two admissible operations is admissible. This is invoked in Proposition 3.4 to conclude that "apply F, then measure O" is an admissible procedure when F and O are individually admissible.

Throughout, δ_{λ} denotes the definite (extremal) state with $w(\delta_{\lambda}, \lambda) = 1$ and $w(\delta_{\lambda}, \mu) = 0$ for $\mu \neq \lambda$.

Cross-configuration structure denotes any aspect of ψ beyond its marginal weight vector $\{w(\psi, \lambda)\}_{\lambda \in \Lambda}$: correlations, phases, joint distributions, or other features that may influence pre-factual dynamics. The existence of such structure is not denied — it is essential to the interference phenomena established in the parent paper. What is asserted is narrower: within the admissible reconstruction algebra, such structure enters solely through its effect on commitment-traceable observables and admissible compositions. Two states that agree on all commitment-traceable observables are identified within the reconstruction algebra as observationally equivalent, not in any deeper ontological sense. Proposition 3.4 will show that admissible dynamics respect this equivalence, so that no admissible operation can produce a recordable distinction between such states. Observational equivalence is therefore not ontological identity; it is the admissible physical equivalence relation relevant to the theory, a status established by that proposition rather than assumed in advance.

Affine weight decomposition. By Proposition 3.6b of the parent paper, the weight function is affine in the composition of pre-factual states. We write $\psi = \psi' + r$ to mean that $w(\psi, \lambda) = w(\psi', \lambda) + w(r, \lambda)$ for all λ, where the right-hand side lies in [0,1] with appropriate normalisation. *Section 4 applies only to affine decompositions that are internally realised within $\mathcal{P}$— not merely formal weight differences in an external linear envelope.* Whenever such a decomposition $\psi = \psi' + r$ is invoked, both ψ' and r are genuine elements of the pre-factual sector, subject to the full dynamics and recombination operations. In what follows, Section 4 is restricted to affine decompositions that are internally realised within $\mathcal{P}$, as permitted by the affine closure structure established in Proposition 3.6b of the parent paper. An example of a decomposition *excluded* by this restriction: if $w(\psi, \lambda_0) < w(\psi', \lambda_0)$ for some λ_0, the formal weight difference $w(\psi, \lambda) - w(\psi', \lambda)$ takes a negative value at λ_0 and does not correspond to any element of $\mathcal{P}$. Section 4 never invokes such decompositions; all residuals r are assumed to have non-negative weights, placing $r \in \mathcal{P}$ as a valid pre-factual state.

The admissibility quotient. By admissibility (A2), physically meaningful distinctions must be traceable to recordable facts. Accordingly, the recombination algebra is understood henceforth on *observational equivalence classes* [ψ] of pre-factual states: two states in the same class are identified as a single algebraic element. Operations and relations on classes are well-defined by the operational congruence principle established in Section 3.4. The null element is the unique class [0] with vanishing marginal weights at all λ. This convention is global and active throughout Sections 3 and 4.

---

3. From Admissibility to Marginal Compositional Consistency

3.1 The Gap in the Original Derivation

The original (A5) stated that the update of w(ψ, λ) under composition with δ_{μ} depends only on w(ψ, λ) and t. The vulnerability: two states ψ and ψ' might have equal marginal weights at every λ yet differ in cross-configuration structure. If that structure influences the marginal update, (A5) fails without violating any constraint the parent paper explicitly imposed. Closing this gap requires a formal principle establishing that admissible dynamics cannot depend on structure finer than observational equivalence.

3.2 Commitment-Traceable Observables: Justification and Scope

Definition 3.1 (Commitment-traceable observable). An observable O on $\mathcal{P}$ is *commitment-traceable* if there exists $f: \Lambda \to \mathbb{R}$ such that

$\langle O \rangle_{\psi} = \sum_{\lambda \in \Lambda} w(\psi, \lambda) \cdot f(\lambda)$ for all $\psi \in \mathcal{P}$.

Remark 3.1a (Justification and ontological scope). The restriction to commitment-traceable observables is not a restriction by fiat. A fact (Definition 2.1 of the parent paper) is a *stable, irreversible record distinguishing one configuration from all others in* Λ. Every admissible observable must be traceable to facts — to configurations recorded by the single-configuration commitment process $\varphi: \mathcal{P} \to \Lambda$. An observable whose expectation depends on cross-configuration structure beyond the marginal weights would require facts recording correlations between configurations *not* committed to in that run. Uncommitted configurations leave no record. Definition 3.1 is therefore a consequence of the definition of a fact.

This does not deny that cross-configuration structure exists or that it has physical consequences. Interference phenomena reveal such structure statistically, across many runs — but each individual run commits to a single configuration, and the statistical distribution across runs is captured entirely by the marginal weights. The admissible physical significance of any pre-factual state is exhausted by its effect on the distribution of committed facts. This quotient does not deny the possible existence of richer pre-factual structure; it asserts only that such structure enters the present reconstruction solely through its effect on commitment-traceable observables and admissible compositions.

Scope note. This argument applies within the single-configuration commitment framework. Theories admitting joint commitment maps $\varphi: \mathcal{P} \to \wp(\Lambda)$ would require a broader treatment.

3.3 Observable Equivalence

Definition 3.2 (Observable equivalence). States $\psi, \psi' \in \mathcal{P}$ are *observationally equivalent* if $\langle O\rangle_{\psi} = \langle O\rangle_{\psi'}$ for every commitment-traceable observable O. We write $\psi \cong \psi'$.

Lemma 3.3 (Characterisation of observable equivalence). $\psi \cong \psi'$ if and only if $w(\psi, \lambda) = w(\psi', \lambda)$ for all $\lambda \in \Lambda$.

Proof. Equal marginal weights give equal expectations for all commitment-traceable observables. Unequal weight at some λ_0 is detected by $f = \mathbb{1}_{\{\lambda_0\}}$. □

3.4 Operational Congruence of Admissible Dynamics

Proposition 3.4 (Operational congruence). Admissibility (A2) requires that observational equivalence is a congruence for all admissible physical operations. That is, for any admissible operation F on the pre-factual sector:

$$\psi \cong \psi' \Longrightarrow F(\psi) \cong F(\psi')$$

Proof. Suppose $F(\psi) \ncong F(\psi')$ for some $\psi \cong \psi'$. By Lemma 3.3, ψ and ψ' agree on all marginal weights and are therefore indistinguishable by any commitment-traceable observable before F is applied. The operation F is a fixed admissible physical process. If $F(\psi) \ncong F(\psi')$, then there exists a commitment-traceable observable O such that $\langle O\rangle_{F(\psi)} \neq \langle O\rangle_{F(\psi')}$. But then the composite procedure "apply F, then measure O" defines an admissible means of distinguishing ψ from ψ' — since the composite of admissible operations is admissible, and its output discriminates the two inputs. This composite procedure would make physically significant an input distinction not grounded in any commitment-traceable content, since $\psi \cong \psi'$. By admissibility (A2), distinctions not traceable to recordable facts carry no physical meaning. Contradiction. □

Remark 3.5. Proposition 3.4 is the formal bridge the MCC derivation requires. It is important to distinguish two levels. *Observational equivalence* (Definition 3.2) is defined purely in terms of commitment-traceable observables at the input level: $\psi \cong \psi'$ means they agree on all such observables before any operation is applied. *Admissible physical equivalence* is the stronger claim that no admissible operation — including composition followed by measurement — can produce a recordable distinction between them. Proposition 3.4 establishes that observational equivalence implies admissible physical equivalence, as a dynamical consequence of admissibility (A2). This extension is the work of the proposition; it should not be read into the phrase "observational equivalence" before that work has been done.

3.5 Deriving Marginal Compositional Consistency

Definition 3.6 (Marginal Compositional Consistency — MCC). A composition rule $\star_t$ satisfies *marginal compositional consistency* if for all $\psi, \psi' \in \mathcal{P}$, all $\mu \in \Lambda$, and all $t \in [0,1]$:

$$\psi \cong \psi' \Longrightarrow \psi \star t\, \delta\{\mu\} \cong \psi' \star t\, \delta\{\mu\}$$

Remark 3.7 (MCC is strictly weaker than (A5)). (A5) requires $w(\psi \star t\, \delta\{\mu\}, \lambda)$ to depend only on $w(\psi, \lambda)$ and t — a pointwise condition excluding dependence on $w(\psi, \mu')$ for any $\mu' \neq \lambda$. MCC requires only that the output depends on the full marginal weight vector and on nothing beyond it. The weight at λ after composition may depend on $w(\psi, \mu')$ for $\mu' \neq \lambda$, provided it depends on nothing finer than the marginal weight vector. (A5) implies MCC; MCC does not imply (A5).

Theorem 3.8 (MCC from admissibility). Under admissibility (A2) and Definition 3.1, any composition rule satisfying conditions (i)–(vi) of Proposition 3.6b of the parent paper satisfies MCC.

Proof. Composition $\star t\, \delta\{\mu\}$ at fixed μ and t is an admissible physical operation. By Proposition 3.4, admissibility requires it to be a congruence for observational equivalence. MCC is precisely that congruence condition for this class of operations. □

Remark 3.9 (Scope). The proof depends on Remark 3.1a and the single-configuration commitment framework.

3.6 Restriction to the Two-Alternative Sector

Remark 3.10 (Why the two-alternative sector is sufficient). The Cauchy functional equation argument operates on states with support on a two-element subset $\{\lambda_1, \lambda_2\} \subseteq \Lambda$. Three observations justify this.

First, the non-classicality conclusion is established by falsifying support-monotonicity in any embedded two-alternative subsystem. A single violation suffices to show the rule is not a convex mixture globally.

Second, the affine reduction (Proposition 3.6b of the parent paper) is proved there for all pairs of pre-factual states. That proposition establishes that a composition rule affine on every two-alternative sector is affine everywhere — so a single sector violation falsifies global affineness.

Third, once one sector violates support-monotonicity, the global theory is non-classical. The interference theorem holds globally from the two-alternative sector result without additional multi-configuration assumptions.

3.7 Recovering Multiplicative Attenuation

Theorem 3.11 (Multiplicative attenuation under MCC). Under MCC (Theorem 3.8) and the semigroup law (v), for extremal states $\psi_1 = \delta_{\lambda_1}$ and $\psi_2 = \delta_{\lambda_2}$ with $\lambda_1 \neq \lambda_2$:

$$w((\psi_1 \star_s \psi_2) \star_t \psi_2, \lambda_1) = w(\psi_1 \star_s \psi_2, \lambda_1) \cdot w(\delta_{\lambda_1} \star_t \psi_2, \lambda_1)$$

Proof. On the two-alternative sector $\{\lambda_1, \lambda_2\}$, the marginal weight vector is fully characterised by $p = w(\theta, \lambda_1) \in [0,1]$. By MCC, the outcome $\theta \star_t \delta_{\lambda_2}$ depends only on the marginal weight vector; on this one-dimensional sector that reduces to dependence on p alone. Write $\varphi_t(p)$ for $w(\theta \star_t \delta_{\lambda_2}, \lambda_1)$ when $w(\theta, \lambda_1) = p$. The semigroup law (v) applied to δ_{λ_1} gives $\varphi_t(\varphi_s(1)) = \varphi_{s+t(1-s)}(1)$. Setting $p(s) = \varphi_s(1)$: $p(s + t(1-s)) = \varphi_t(p(s))$. By MCC, any state with $w(\theta, \lambda_1) = p(s)$ yields the same φ_t, so for $\theta = \psi_1 \star_s \psi_2$:

$$w((\psi_1 \star_s \psi_2) \star_t \psi_2, \lambda_1) = \varphi_t(p(s)) = p(s + t(1-s)) = p(s) \cdot p(t)$$

where the final equality is the multiplicative Cauchy equation from parent paper Lemma 3.6b'. □

Remark 3.12. On the two-alternative sector the marginal weight vector is one-dimensional, so MCC and (A5) coincide on this sector. MCC is strictly weaker globally; they are equivalent where the Cauchy argument operates.

---

4. From Partial Cancellation to Algebraic Closure

Level distinction. Condition (S2) — universal additive inverses — is a statement about **A**, the amplitude algebra formed on admissibility classes of $\mathcal{P}$, not a claim that elements of $\mathcal{P}$ have probability-weight inverses within $\mathcal{P}$. The contradiction in Theorem 4.10 is driven by $\mathcal{P}$-level weight considerations — IAC gives $\chi \in \mathcal{P}$ with $r + \chi \cong 0$, non-negativity of w forces $w(r, \lambda) = 0$ for all λ — with the final lift to algebraic nullity supplied by faithfulness. The conclusion at the A level then follows: since no non-null irreducible residual survives in $\mathcal{P}$, the algebra A on admissibility classes is closed under additive inversion, establishing S2 as a property of A.

4.1 The Gap in the Original Derivation

The parent paper's Theorem 3.7 established pointwise cancellability: for any non-classical regular composition rule, some compositions drive the weight at a given configuration to zero. Earlier versions then attempted to extend this to universal additive inverses via a reversible-closure condition (RC) on all physically meaningful reversible transformations. This was criticised as too broad.

The present version replaces RC with *internal admissible closure* (IAC) — a precisely stated condition covering only observationally active, internally realised affine sub-contributions. The contribution is a reduction: the old broad closure requirement is replaced by this narrower principle, from which (S2) follows. Whether IAC itself follows from more primitive premises is the primary open question.

A further gap in earlier versions concerned the step from observational nullity ($r + \chi \cong 0$) to algebraic nullity ($r + \chi = 0$). This is bridged by the Faithfulness Lemma (4.4), proved as a preliminary before the main chain, using the admissibility quotient declared globally in Section 2.

4.2 Irreducible Residuals

Definition 4.1 (Pointwise cancellability). The pre-factual sector is *pointwise cancellable* if for every $\psi \in \mathcal{P}$ and every $\lambda \in \Lambda$ 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 follows from Theorem 3.7 of the parent paper for any non-classical regular composition rule.

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

- **(i)** $r \neq 0$ (r is non-null),
- **(ii)** $\psi = \psi^0 + r$ in the affine weight decomposition of Section 2, internally realised in $\mathcal{P}$, for some $\psi^0 \in \mathcal{P}$,
- **(iii)** there exists no $\chi \in \mathcal{P}$ with $r + \chi \cong 0$ (r has no observational cancellation partner in $\mathcal{P}$).

Non-nullity (i) guarantees $w(r, \lambda_0) > 0$ for at least one $\lambda_0 \in \Lambda$, since weights are non-negative. No separate activity condition is needed.

4.3 Internal Admissible Closure

Definition 4.3 (Internal admissible closure — IAC). The pre-factual sector satisfies *internal admissible closure* if: for any internally realised affine sub-contribution $r \in \mathcal{P}$ with $\psi = \psi' + r$, if there exists an admissible composition sequence — including, in particular, the identity case $t = 0$ — after which r induces a difference in commitment-traceable observables between ψ and ψ', then there exists $\chi \in \mathcal{P}$ such that $r + \chi \cong 0$.

Note on the identity case. The composition family includes identity at $t = 0$ (see Section 2). A non-null r with $w(r, \lambda_0) > 0$ induces a difference in commitment-traceable observables between $\psi = \psi' + r$ and ψ' at $t = 0$, without requiring any non-trivial composition. IAC is therefore triggered by the initial weights alone for any non-null internally realised r.

Remark 4.4 (Scope of IAC). IAC asserts only that internally realised affine sub-contributions with commitment-detectable physical content have observational recombination partners. It does not assert global operational completeness, that all reversible transformations are realisable, or that every state has an inverse. Its scope is narrow: it applies only to the specific class of non-null, internally realised affine sub-contributions.

Remark 4.5 (Physical motivation). If an affine sub-contribution r of a pre-factual state ψ has genuine commitment-traceable physical content, the transition from ψ to $\psi^0 = \psi - r$ is a physically real transition within the pre-factual domain. A theory whose recombination algebra cannot express this transition would contain physically real pre-factual transitions not representable within its stated dynamics. IAC is the condition that the algebra is complete with respect to this specific class of transitions. Whether this follows from recombinability (R2) and admissibility (A2) alone is open question 1 in Section 7.

4.4 Preliminary Lemma: Faithfulness of the Null Element

Lemma 4.6 (Faithfulness of the null element). If $\eta \in \mathcal{P}$ satisfies $\eta \cong 0$, then $\eta = 0$ as an algebraic element.

Proof. Under the admissibility quotient declared in Section 2, the algebra operates on observational equivalence classes. The null element is the unique class [0] with vanishing marginals at all λ. $\eta \cong 0$ places η in that class. Therefore $[\eta] = [0]$, i.e. $\eta = 0$ as an algebraic element. □

Remark 4.6a. This lemma is immediate given the admissibility quotient stated globally in Section 2. No local redefinition is required here; the quotient was the operating convention from the start.

4.5 Lemma: Irreducible Residuals Induce Admissible Internal Asymmetry

Lemma 4.7. Let r be an irreducible residual of ψ with decomposition $\psi = \psi^0 + r$. Then $\psi \ncong \psi^0$.

Proof. Since $r \neq 0$ and $w(r, \lambda) \geq 0$ for all λ, 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 Lemma 3.3, $\psi \ncong \psi^0$. □

4.6 Lemma: Admissible Internal Asymmetry Is Recombination-Realisable under IAC

Lemma 4.8. Under IAC (Definition 4.3), if r is an irreducible residual of ψ, then there exists $\chi \in \mathcal{P}$ with $r + \chi \cong 0$.

Proof. By Lemma 4.7, $w(\psi, \lambda_0) \neq w(\psi^0, \lambda_0)$. At $t = 0$ (the identity composition, admissible by the family definition in Section 2), r induces a commitment-traceable difference between ψ and ψ^0. The condition of Definition 4.3 is met. By IAC, there exists $\chi \in \mathcal{P}$ with $r + \chi \cong 0$. □

4.7 Lemma: Observational Nullity Implies Marginal Weight Cancellation

Lemma 4.9. If there exists $\chi \in \mathcal{P}$ with $r + \chi \cong 0$, then $w(r, \lambda) + w(\chi, \lambda) = 0$ for all $\lambda \in \Lambda$.

Proof. $r + \chi \cong 0$ means $r + \chi$ is in the null admissibility class [0], which by definition (Section 2) has vanishing marginals at all λ, i.e. $w(0, \lambda) = 0$ for all λ. By Lemma 3.3, $r + \chi \cong 0$ implies $w(r + \chi, \lambda) = w(0, \lambda) = 0$ for all λ. By the affine weight decomposition, $w(r + \chi, \lambda) = w(r, \lambda) + w(\chi, \lambda)$. Hence $w(r, \lambda) + w(\chi, \lambda) = 0$ for all λ. □

4.8 Theorem: No Irreducible Residuals

Theorem 4.10 (Cancellability closure). Suppose the pre-factual sector satisfies pointwise cancellability (PC), continuity (R1), recombinability (R2), exchange symmetry (vi), admissibility (A2), and internal admissible closure (IAC). Then no irreducible residuals exist, and the pre-factual sector is globally cancellable.

Proof. The assumptions (R1), (R2), and (vi) enter through the parent-paper results establishing the surrounding recombination framework and pointwise cancellability. The contradiction step below uses only the consequences isolated in Lemmas 4.6–4.9 together with IAC. Assume for contradiction that ψ has an irreducible residual r — in particular $r \neq 0$ (condition (i)). By Lemma 4.7, $\psi \ncong \psi^0$. By Lemma 4.8, IAC yields $\chi \in \mathcal{P}$ with $r + \chi \cong 0$. By Lemma 4.9, $w(r, \lambda) + w(\chi, \lambda) = 0$ for all λ. Since both $w(r, \lambda) \geq 0$ and $w(\chi, \lambda) \geq 0$, this forces $w(r, \lambda) = 0$ for all λ — that is, $r \cong 0$. By Lemma 4.6 (faithfulness), $r = 0$ as an algebraic element. This contradicts condition (i) of Definition 4.2. □

Corollary 4.11 (Universal additive inverses — S2). Under the conditions of Theorem 4.10, condition (S2) holds in the amplitude algebra A: every admissibility class $[\psi] \in A$ admits an additive inverse $[-\psi]$ such that $[\psi] + [-\psi] = [0]$. Since A is the algebra on admissibility classes of $\mathcal{P}$, the absence of non-null irreducible residuals in $\mathcal{P}$ implies the existence of such an inverse class for every $[\psi] \in A$.

Proof. Let $r \in \mathcal{P}$ be any non-null element. Condition (i) of Definition 4.2 holds by assumption. Condition (ii) holds trivially: taking $\psi = r$ and $\psi^0 = 0$ gives the internally realised decomposition $r = 0 + r$ with both $r \in \mathcal{P}$ and $0 \in \mathcal{P}$. With both (i) and (ii) satisfied, Theorem 4.10 implies r is not an irreducible residual, so condition (iii) must fail: there exists $\chi \in \mathcal{P}$ such that $r + \chi \cong 0$. By Lemma 4.6 (faithfulness), $r + \chi = 0$ as an algebraic element, so $[r] + [\chi] = [0]$ in A. Hence $[\chi]$ is the additive inverse of [r] in A. Since r was arbitrary, every non-null class in A has an additive inverse; the null class [0] is its own inverse. □

4.9 Honest Scope

Logical dependencies of the proof chain:

- Lemma 4.6 (preliminary): $\eta \cong 0 \rightarrow \eta = 0$. Uses the admissibility quotient of Section 2 (declared globally).
- Lemma 4.7: irreducibility → observational distinguishability. Uses non-nullity and non-negativity of weights only.
- Lemma 4.8: observational distinguishability → observational cancellation partner $\chi \in \mathcal{P}$ with $r + \chi \cong 0$. Invokes IAC here and only here.
- Lemma 4.9: $r + \chi \cong 0 \rightarrow w(r,\lambda) + w(\chi,\lambda) = 0$ for all λ. Uses affine structure, Lemma 3.3, and null state definition.
- Theorem 4.10: non-negativity forces $w(r,\lambda) = 0$; faithfulness (Lemma 4.6) gives $r = 0$; contradicts condition (i).

IAC enters at one step — Lemma 4.8. The faithfulness lemma is proved before the chain and used transparently at the closing step. The observational-to-algebraic nullity gap present in all earlier versions is closed.

On the narrowness of the reduction. The note on the identity case in Definition 4.3 establishes that IAC is triggered by the initial weights alone for every non-null internally realised element. IAC is therefore closely adjacent to S2 at the observational level: every non-null internally realised element has an observational cancellation partner in $\mathcal{P}$. The substantive logical move is the combination of Lemma 4.9 with faithfulness: non-negativity first forces $r \cong 0$, and Lemma 4.6 then lifts this observational nullity to algebraic nullity $r = 0$. The genuine contribution of this section is replacing the old broad reversible-closure condition with a precisely stated, $\mathcal{P}$-level principle that admits a transparent proof chain rather than a vague appeal to closure under all reversible transformations.

The residual openness. IAC is not derived from (A1) and (A2) alone. The strongest accurate claim is: we have reduced the old broad closure requirement to a narrower condition restricted to the specific class needed by the proof, and derived (S2) from it.

---

5. Revised Algebraic Structure Theorem

Theorem 5.1 (Revised division algebra structure). Let a physical theory satisfy (A1), (A2), (A4), finite-dimensionality over $\mathbb{R}$, and IAC; and let composition satisfy (v), (vi), (R1), (R2). Then:

(a) Operational congruence (Proposition 3.4) holds, MCC follows (Theorem 3.8), and multiplicative attenuation is recovered (Theorem 3.11), yielding the Cauchy equation and structural interference.

(b) Pointwise cancellability extends to global cancellability under IAC (Theorem 4.10), establishing (S2) (Corollary 4.11) and abelian group structure on A.

(c) With (A4) and the abelian group structure established in (b), the division-algebra conclusion of Theorem 5.2 of the parent paper applies: A carries the structure of a finite-dimensional division algebra over $\mathbb{R}$. By the Frobenius theorem the only candidates are $\mathbb{R}$, $\mathbb{C}$, $\mathbb{H}$.

---

6. Impact on the Bridge Theorem

Theorem 6.1 (Revised Bridge Theorem). Under the conditions of Theorem 5.1, all results of the parent paper's Bridge Theorem hold, with the following changes:

Previous assumption	Parent paper status	Status here
(A5) Locality of composition	Primitive — open question	**Derived** from (A2) via Proposition 3.4 + Theorem 3.8
(S2) Universal additive inverses	Primitive — motivated by (A2)	**Derived under IAC** via Lemmas 4.6–4.9 + Theorem 4.10
(IAC) Internal admissible closure	Not stated — implicit in RC	**Made explicit** — narrower than RC — open question 1

The remaining results follow as before.

Remark 6.2 (Net logical gain). (A5) is derived from admissibility via a named intermediate proposition, with the two levels of equivalence — observational and admissible physical — clearly distinguished. (S2) is derived from IAC via a four-lemma chain with a preliminary faithfulness result; the contribution is a reduction to a more sharply delimited condition covering only the specific class of internally realised, observationally active affine sub-contributions. Every load-bearing step is explicitly labelled. No algebraic postulate is left invisible, unmotivated, or unnamed.

---

7. Remaining Assumptions and Open Questions

1. Internal admissible closure (IAC). Whether Definition 4.3 follows from (A1), (A2), and (R2) alone is the primary open question. Remarks 4.4 and 4.5 give physical motivation but not a derivation.

1a. Admissibility closed under composition (A2c). Proposition 3.4 invokes the principle that the composition of two admissible operations is admissible. This is stated as an explicit condition in Section 2 but is not derived from (A1)–(A2) alone. Whether it follows from the definition of admissibility or requires independent postulation is an open question.

2. Metric homogeneity (A3). Deriving the manifold and metric structure of Λ from fact-production constraints remains open.

3. Commutativity (C5). Needed to reach a commutative field and Galois invariance. Not derivable from (A1)–(A4).

4. Finite-dimensionality. Required for the Frobenius theorem. Physically motivated; not derived.

5. Semigroup law (v) and exchange symmetry (vi). Admissibility-based motivations are in the parent paper. Full derivations from (A1)–(A2) alone are not given.

6. Imported distinguishability geometry results. Exclusion of $\mathbb{R}$ and $\mathbb{H}$ remain imported from the companion paper on distinguishability geometry.

---

8. Conclusion

We have addressed two structural gaps in the VERSF Bridge Theorem.

(A5) is derived. The derivation introduces an explicit intermediate step — the *operational congruence principle* (Proposition 3.4) — establishing that admissible dynamics must act on observational equivalence classes, because any operation mapping equivalent inputs to distinguishable outputs would render physically meaningful a distinction not grounded in commitment-traceable facts. MCC (Theorem 3.8) follows immediately. The key bridge, previously implicit and compressed within a single sentence, is now a named formal proposition with its own proof.

The ontological scope of commitment-traceable observables is clarified: richer cross-configuration structure is not denied. Such structure underlies interference phenomena and exists in the pre-factual sector. What is asserted is that within the admissible reconstruction algebra, such structure enters solely through its effect on commitment-traceable observables and admissible compositions. Observational equivalence is not ontological identity; Proposition 3.4 shows that it defines the admissible physical equivalence relation relevant to the reconstruction.

(S2) is derived under IAC. The contribution is a reduction, not a full derivation from fact-production alone. We show that the broad reversible-closure condition (RC) of earlier versions can be replaced by IAC — a precisely stated, narrower principle covering only non-null, internally realised affine sub-contributions. The proof chain runs through four lemmas, preceded by a preliminary Faithfulness Lemma (4.6) proved before the chain using the admissibility quotient declared globally in Section 2. IAC enters at one step (Lemma 4.8), yielding a cancellation partner $\chi \in \mathcal{P}$ with $r + \chi \cong 0$; non-negativity of weights then forces $w(r, \lambda) = 0$ for all λ, so $r \cong 0$; faithfulness (Lemma 4.6) gives $r = 0$, contradicting condition (i) of Definition 4.2.

No algebraic postulate about the pre-factual sector is left invisible, unmotivated, or unnamed. The remaining nontrivial structural burden is therefore concentrated in IAC and the imported geometric assumptions, rather than dispersed across hidden algebraic postulates.