Why a Fact-Producing Universe Must Satisfy Interference, Isotropy, and Representational Invariance

Keith Taylor
VERSF Theoretical Physics Program

For the General Reader

Quantum mechanics is strange. Possibilities interfere with each other. Outcomes are irreducibly probabilistic. Complex numbers appear at the foundation of physical law. Symmetry is not a feature but a requirement. Why?

The usual answer is: that is just how the universe is. The equations work; the predictions match experiment; no further explanation is offered.

This paper offers a different kind of answer — not by assuming quantum mechanics and checking that it fits, but by asking what any universe must look like if it is to produce definite, stable facts at all. A universe in which nothing ever really *happens* — in which no measurement ever settles, no particle ever lands, no event ever becomes irreversible — is not a universe at all. The requirement that facts exist turns out to be surprisingly powerful.

We show that from this single requirement, together with a small number of consistency conditions, the key structural features of quantum mechanics follow by a constrained chain of argument. Possibilities must be able to cancel each other before a fact is formed: that is interference. Configurations that are equally distinguishable must be treated equally: that is symmetry. The mathematical language describing unresolved possibilities must not smuggle in distinctions that no fact can ever record: that is invariance. And the algebra of that language must be closed, invertible, and consistent: that narrows the candidates to a small class of mathematical structures — of which complex numbers are the unique member satisfying all remaining constraints.

We are careful about what we claim. Some steps in the derivation are complete. Others depend on results established in prior work, which we cite explicitly. And we identify honestly the assumptions we cannot yet derive — the geometry of the configuration space, the commutativity of amplitude multiplication, the dimension of the resulting Hilbert space — leaving them as open questions rather than absorbing them silently.

The conclusion is not that quantum mechanics is merely consistent with a fact-producing universe. It is that quantum mechanics is the minimal mathematical structure such a universe can have.

Abstract

We address a foundational question: what structural constraints are imposed on a physical theory by the requirement that it produce stable, irreversible facts from unresolved alternatives? Previous work has shown that complex Hilbert space arises uniquely from distinguishability geometry together with interference, isotropy, and Galois invariance. Here we derive these structural principles from more primitive requirements, and address gaps left open in earlier treatments.

We establish four main structural steps. First, any theory supporting a genuinely non-classical reversible pre-commitment sector cannot be equivalent to a classical stochastic model; this necessitates non-classical composition of unresolved alternatives. Second, with a regularity condition on that composition, non-classical structure implies phase-sensitive cancellation, i.e. interference. Third, admissibility of physical distinctions implies permutation symmetry over equally distinguishable configurations; under metric homogeneity this extends to continuous isotropy. Fourth, associative, bilinear composition together with invertibility forces the amplitude domain to carry *division algebra* structure; under an additional commutativity assumption this reduces to a field, and representational invariance then yields Galois invariance.

We explicitly address the selection among $\mathbb{R}$, $\mathbb{C}$, and $\mathbb{H}$, identifying the precise constraints that exclude each non-complex candidate, and we distinguish results established in this paper from those imported from prior work on distinguishability geometry. Together, the results show that the structural assumptions required for the emergence of complex Hilbert space are not arbitrary, but follow from the requirement of consistent fact-production in a non-classical universe.

Table of Contents

- For the General Reader
- Abstract

1. Introduction
2. Fact-Production and the Pre-Factual Sector
3. Necessity of Non-Classical Composition and Interference
   - 3.1 Classical Composition is Insufficient
   - 3.2 Non-Classical Composition Implies Interference (Proposition 3.6b with semigroup law (v) and exchange symmetry (vi); Theorem 3.7)
4. Necessity of Isotropy
   - 4.1 Permutation Symmetry
   - 4.2 Extension to Continuous Isotropy
5. Division Algebra Structure and the Origin of Fields
   - 5.1 Amplitudes Without Born-Rule Presupposition (Lemma 5.1a; Theorem 5.2)

6. Necessity of Representational Invariance and Galois Invariance
7. Selection of Complex Amplitudes
    - 7.1 Ruling Out $\mathbb{R}$: Failure of Continuous Isotropy for Composite Systems
    - 7.2 Ruling Out $\mathbb{H}$: Failure of Representational Invariance for Composite Systems
    - 7.3 Complex Amplitudes Satisfy All Constraints
8. The Bridge Theorem
9. Responses to Anticipated Objections
10. Discussion
11. Conclusion

- Appendix: Logical Dependency Map

---

1. Introduction

Quantum mechanics is characterised by a specific mathematical structure: complex amplitudes, interference, symmetry, and Hilbert space. Reconstruction programmes have derived this structure from operational or information-theoretic axioms. In previous work within the VERSF framework, we showed that complex Hilbert space is uniquely selected by the joint requirements of interference, isotropy, and Galois invariance.

The present paper addresses the status of those requirements themselves. We begin from a single minimal premise:

A physical universe must be capable of producing stable, irreversible facts from unresolved alternatives.

We show that this premise, together with an admissibility constraint on physical distinctions and regularity conditions on the compositional structure, forces the structural principles underlying quantum theory. The derivation proceeds in five steps:

1. **Non-classical composition** (Section 3): required to prevent collapse of the pre-factual sector into classical ignorance.
2. **Interference** (Section 3): follows from non-classical composition under a continuity and recombination condition.
3. **Isotropy** (Section 4): permutation symmetry follows from admissibility; continuous isotropy follows under metric homogeneity of the configuration space.
4. **Division algebra and field structure** (Section 5): associative, bilinear, invertible composition forces the amplitude domain to carry division algebra structure; commutativity, when it holds, reduces this to a field.
5. **Representational invariance / Galois invariance** (Section 6): admissibility forces predictions to be invariant under norm-preserving automorphisms; when the amplitude domain is a field this is Galois invariance.

Section 7 addresses the selection of $\mathbb{C}$ over $\mathbb{R}$ and $\mathbb{H}$, identifying which constraints each fails and where prior work is relied upon. Section 8 states the Bridge Theorem with explicit labelling of what is new and what is imported.

---

2. Fact-Production and the Pre-Factual Sector

We assume the following minimal structure:

- A set of **distinguishable configurations** Λ, equipped with a distinguishability relation δ: $\Lambda \times \Lambda \rightarrow [0,1]$ where $\delta(\lambda_1, \lambda_2) = 0$ if and only if $\lambda_1 = \lambda_2$.
- A **reversible pre-commitment sector**: a set of pre-factual states $\mathcal{P}$ with reversible dynamics.
- An **irreversible commitment process** φ: $\mathcal{P} \rightarrow \Lambda$ producing facts.
- An **admissibility constraint**: only distinctions recordable as facts carry physical meaning.

Definition 2.1 (Fact). A fact is a stable, irreversible record distinguishing one configuration from all others in Λ.

Definition 2.2 (Pre-factual state). A pre-factual state $\psi \in \mathcal{P}$ is one in which multiple alternatives $\lambda_i \in \Lambda$ remain jointly unresolved.

Definition 2.3 (Nontriviality). The pre-factual sector is *nontrivial* if no pre-factual state $\psi \in \mathcal{P}$ is operationally equivalent to a probability distribution over Λ.

The requirement of nontriviality asserts that unresolved alternatives are genuinely distinct from ignorance of a definite outcome. It is the critical structural demand from which the subsequent arguments proceed.

---

3. Necessity of Non-Classical Composition and Interference

We proceed in two stages: first showing that classical composition is insufficient (Theorem 3.4), then showing that under a regularity condition the required non-classical structure implies interference (Theorem 3.7).

3.1 Classical Composition is Insufficient

Definition 3.1 (Classical stochastic model). A theory is *classical stochastic* if pre-factual states are probability distributions over Λ, dynamics preserves convex structure, and commitment samples from that distribution.

Lemma 3.2 (Classical reducibility). In a classical stochastic theory, every pre-factual state is operationally equivalent to a mixture over definite outcomes.

Proof. Let $p: \Lambda \to [0,1]$ be a pre-factual state with $\Sigma\_\lambda\ p(\lambda) = 1$. For any admissible observable $O: \Lambda \to \mathbb{R}$, the expected outcome is $\langle O \rangle = \Sigma\_\lambda\ p(\lambda)\ O(\lambda)$. This is identical to the expectation under a preparation that selects configuration λ with probability $p(\lambda)$ and suppresses knowledge of the selection. No sequence of admissible operations can distinguish the pre-factual state from this mixture, since all observable outputs are convex combinations of values on Λ in both cases. The state is therefore operationally equivalent to classical ignorance. □

Lemma 3.3 (Pre-factual collapse). In a classical stochastic theory, the distinction between pre-factual and post-commitment states is operationally empty.

Proof. By Lemma 3.2, every pre-factual state is a mixture over Λ. Every post-commitment state is a point distribution on Λ. The commitment process φ merely sharpens the distribution; it does not mediate a transition from a qualitatively distinct domain. No fact can record the difference between "genuinely unresolved" and "resolved but unknown." □

Theorem 3.4 (Non-classical composition necessity). Any theory supporting a nontrivial pre-factual sector must admit a composition rule for alternatives that is not reducible to convex combination.

Proof. Suppose all composition of alternatives proceeds by convex combination. Then every pre-factual state is a convex mixture over Λ. By Lemma 3.2, every such state is operationally equivalent to a classical mixture. This violates nontriviality (Definition 2.3). Therefore the composition rule cannot be convex combination. □

Remark 3.5. Theorem 3.4 establishes that some non-classical composition is required. It does not yet specify *which* non-classical structure. In particular, it does not rule out nonlinear or exotic composition rules. The next step derives interference specifically, under a regularity condition.

3.2 Non-Classical Composition Implies Interference

To proceed from non-classical composition to interference, we introduce a regularity condition.

Definition 3.6 (Regularity). The composition rule for alternatives is *regular* if:

- **(R1) Continuity**: the composition of two alternatives varies continuously with a compositional parameter (e.g. a relative weighting or phase),
- **(R2) Recombinability**: alternatives that have been jointly evolved can be recombined into a single pre-factual state,
- **(R3) Reversibility**: the composition admits an inverse under the dynamics.

Definition 3.6a (Support-monotonicity). A composition rule $\star\_t$ (parametrised by $t \in [0,1]$) is *support-monotone* if for every configuration $\lambda \in \Lambda$ and every pair of pre-factual states ψ_1, ψ_2:

$w(\psi_1 \star\_t\ \psi_2, \lambda) \geq 0$ and $w(\psi_1 \star\_t\ \psi_2, \lambda) \geq \min(w(\psi_1, \lambda), w(\psi_2, \lambda))$

where $w(\psi, \lambda)$ denotes the weight assigned to configuration λ by state ψ. That is, composition never diminishes the weight at any configuration below the lower of the two components.

Proposition 3.6b (Affine reduction under monotone invertible composition). Let $\star\_t$ be a continuous one-parameter family of composition rules on pre-factual states satisfying:

- **(i) Closure**: the composition of two pre-factual states is a pre-factual state,
- **(ii) Support-monotonicity** (Definition 3.6a),
- **(iii) Invertibility**: for each $t \in (0,1]$, the map $\psi_1 \mapsto \psi_1 \star\_t\ \psi_2$ is invertible,
- **(iv) Preservation of extremal states**: definite configurations (point distributions) compose to definite configurations when $t \in \{0,1\}$,
- **(v) Sequential mixing law** (semigroup property): for all $s, t \in [0,1]$ and all pre-factual states ψ_1, ψ_2: $(\psi_1 \star\_s\ \psi_2) \star t\ \psi_2 = \psi_1 \star\{s+t(1-s)\}\ \psi_2$ This says composing ψ_1 with ψ_2 in two sequential steps — first to degree s, then to degree t — is equivalent to a single composition to degree $s + t(1-s) = 1-(1-s)(1-t)$,
- **(vi) Exchange symmetry**: $\psi_1 \star t\ \psi_2$ *and* $\psi_2 \star\{1-t\}\ \psi_1$ produce the same observable state for all $t \in [0,1]$. This reflects the principle that a "t-to-(1-t) mixture" is the same regardless of which state is labelled first.

Then $\star\_t$ is affine on the convex hull of alternatives — that is, $\psi_1 \star\_t\ \psi_2 = (1-t)\psi_1 + t\ \psi_2$ for all $t \in [0,1]$.

We first establish a lemma that makes the Cauchy step in the proof rigorous.

Lemma 3.6b' (Local attenuation). Under conditions (i), (iv), (v), and (vi) of Proposition 3.6b, for extremal states $\psi_1 = \delta\_\{\lambda_1\}$ and $\psi_2 = \delta\_\{\lambda_2\}$, the weight at λ_1 evolves multiplicatively under sequential composition with ψ_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. By condition (iv), $\psi_2 = \delta\_\{\lambda_2\}$ is an extremal state with $w(\psi_2, \lambda_1) = 0$. The semigroup law (v) gives:

$(\psi_1 \star\_s\ \psi_2) \star t\ \psi_2 = \psi_1 \star\{s+t(1-s)\}\ \psi_2$

The left side involves composing the intermediate state $\psi_1 \star s\ \psi_2$ *with* ψ_2 *to degree t. By condition (v) applied to the pair (*$\delta\{\lambda_1\}, \psi_2$), the operation "compose with ψ_2 to degree t" acts on any state θ as $\theta \star t\ \psi_2$. *Because* $\psi_2 = \delta\{\lambda_2\}$ assigns zero weight to λ_1, and by exchange symmetry (vi) the weight at λ_1 cannot be increased above the weight λ_1 carries in the first argument, the weight at λ_1 after the t-composition depends only on: (a) the weight at λ_1 in the first argument θ, and (b) the degree-t attenuation factor determined by the pair $(\delta\_\{\lambda_1\}, \psi_2)$ alone.

We invoke the following assumption:

(A5) Locality of composition: the update of weight at any configuration λ under composition depends only on the prior weight at λ and the composition parameter t, not on the detailed distribution of weight over other configurations.

This locality condition is implicit in the semigroup law (v) when applied uniformly across configurations: if the weight at λ_1 could depend on the distribution at λ_2 during the t-step, the semigroup law could not hold for all first arguments simultaneously without introducing cross-configuration coupling that would conflict with the independence of configurations as distinct distinguishable outcomes. It is physically natural — configuration λ_1 does not "know" about the weight at λ_2 except through the global composition parameter t.

Clarification on scope. (A5) does not assert that configurations evolve independently at the level of the full pre-factual state. It asserts only that, under composition with a definite alternative, the marginal update rule for a configuration depends on its prior weight at that configuration and the composition parameter. Global constraints such as normalisation and consistency across configurations remain fully enforced. Thus (A5) is a locality condition on the *functional form of the update rule*, not a claim of dynamical independence between configurations. Quantum-mechanical global structure — normalisation, phase relations, entanglement — is not excluded by (A5); it operates at a level above the marginal weight update that (A5) governs. In particular, (A5) excludes update rules in which the weight at λ depends on correlations between λ and other configurations that are not themselves part of the composition parameter, as such rules would violate the semigroup property (v) when applied uniformly across all configurations.

Remark on compatibility with quantum mechanics. A reviewer might object that quantum probability amplitudes do not update multiplicatively configuration by configuration — that is precisely what interference breaks. This objection, while well-motivated, conflates two levels of description. (A5) governs the update of *marginal weights under composition with a definite alternative* — it is a statement about how the probability of a configuration changes when one mixes the pre-factual state toward a definite outcome. In quantum mechanics, this marginal update is governed by the Born rule applied to the reduced state, and does satisfy the multiplicative form. To verify: let ρ be a quantum state and let $p(\lambda_1) = \mathrm{Tr}(\rho\, \Pi_{\lambda_1})$ be the Born-rule probability of configuration λ_1. A degree-t composition toward the definite state δ_{λ_2} is represented at the density-matrix level by the convex combination $\rho_t = (1-t)\rho + t|\lambda_2\rangle\langle\lambda_2|$. The updated marginal probability at λ_1 is then:

$$p'(\lambda_1) = \mathrm{Tr}(\rho_t\, \Pi_{\lambda_1}) = (1-t)\, \mathrm{Tr}(\rho\, \Pi_{\lambda_1}) + t\, \mathrm{Tr}(|\lambda_2\rangle\langle\lambda_2|\, \Pi_{\lambda_1}) = (1-t)\, p(\lambda_1) + t \cdot 0 = (1-t)\, p(\lambda_1)$$

where the last step uses $\Pi_{\lambda_1} \Pi_{\lambda_2} = 0$ (orthogonality of distinct configurations). This is precisely $p'(\lambda_1) = p(\lambda_1) \cdot (1-t)$, confirming the multiplicative attenuation form that (A5) requires. The interference between paths — the distinctively quantum feature — lives entirely in the off-diagonal terms of ρ, which affect $p(\lambda_1)$ through their contribution to $\mathrm{Tr}(\rho\, \Pi_{\lambda_1})$ but do not alter the functional form of the marginal update under composition with a definite alternative. (A5) is therefore satisfied in quantum mechanics at the weight level, as claimed.

Under (A5), define the attenuation function $a(t) = w(\delta_{\lambda_1} \star_t \psi_2, \lambda_1)$, *which by definition equals* $p(t)$ *(the weight at* λ_1 *when starting from* δ_{λ_1}). Then for any state θ with support concentrated on $\{\lambda_1, \lambda_2\}$ (as $\psi_1 \star_s \psi_2$ is, by closure and support-monotonicity), (A5) gives directly:

$$w(\theta \star_t \psi_2, \lambda_1) = w(\theta, \lambda_1) \cdot a(t) = w(\theta, \lambda_1) \cdot p(t)$$

The multiplicative factorisation follows: $w((\psi_1 \star_s \psi_2) \star_t \psi_2, \lambda_1) = w(\psi_1 \star_s \psi_2, \lambda_1) \cdot p(t) = p(s) \cdot p(t)$. □

Proof of Proposition 3.6b. We work with extremal states $\psi_1 = \delta_{\lambda_1}$, $\psi_2 = \delta_{\lambda_2}$. Define $p(t) = w(\psi_1 \star_t \psi_2, \lambda_1)$, the weight assigned to λ_1 in the composed state. Boundary conditions give $p(0) = 1$ and $p(1) = 0$; support-monotonicity gives $p(t) \in [0,1]$ for all t; continuity follows from (R1).

Step 1: Cauchy equation from Lemma 3.6b' and the sequential mixing law (v). By Lemma 3.6b', the weight at λ_1 under sequential composition factorises as:

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

Combined with the semigroup law (v), which gives $(\psi_1 \star_s \psi_2) \star_t \psi_2 = \psi_1 \star_{s+t(1-s)} \psi_2$, *and hence* $w(\psi_1 \star_{s+t(1-s)} \psi_2, \lambda_1) = p(s+t(1-s))$, we obtain:

$$p(s + t(1-s)) = p(s) \cdot p(t)$$

Substituting $u = 1-s$, $v = 1-t$ (so $s + t(1-s) = 1-(1-s)(1-t) = 1-uv$), and defining $q(u) = p(1-u)$:

$$q(uv) = q(u) \cdot q(v) \text{ for all } u, v \in [0,1]$$

with $q(0) = p(1) = 0$, $q(1) = p(0) = 1$. This is the **multiplicative Cauchy equation**. The only continuous solutions with $q: [0,1] \to [0,1]$ monotone are $q(u) = u^c$ for $c > 0$. Therefore $p(t) = (1-t)^c$ for some $c > 0$.

Step 2: Pinning c = 1 via exchange symmetry (vi). By condition (vi), $\psi_1 \star_t \psi_2$ *and* $\psi_2 \star_{1-t} \psi_1$ give the same state. Applying Step 1 to the pair (ψ_2, ψ_1), the weight at λ_1 in $\psi_2 \star_{1-t} \psi_1$ equals $1 - \tilde{p}(1-t)$ where $\tilde{p}(t) = w(\psi_2 \star_t \psi_1, \lambda_2) = (1-t)^c$ by the same Cauchy argument. So:

$$w(\psi_2 \star_{1-t} \psi_1, \lambda_1) = 1 - (1-(1-t))^c = 1 - t^c$$

For conditions (vi) to hold, we need this to equal $p(t) = (1-t)^c$ for all $t \in [0,1]$:

$$(1-t)^c = 1 - t^c \text{ for all } t \in [0,1]$$

Setting $t = 1/2$: $2\cdot(1/2)^c = 1$, so $(1/2)^{c-1} = 1$, hence $\mathbf{c = 1}$.

With $c = 1$: $p(t) = 1-t$. Therefore $w(\psi_1 \star_t \psi_2, \lambda_1) = 1-t$ and $w(\psi_1 \star_t \psi_2, \lambda_2) = t$, which is precisely affine mixing. Extending to general states by linearity of the weight function in the components, $\psi_1 \star_t \psi_2 = (1-t)\psi_1 + t\psi_2$. □

Remark 3.6c (Status of conditions (v) and (vi), and scope of Theorem 3.7). Condition (v) is the natural "no memory of intermediate steps" property: the compound degree depends only on s and t, not on the intermediate path. Condition (vi) — exchange symmetry — merits further comment. It can be motivated by admissibility: if labelling ψ_1 "first" versus "second" could affect observable predictions, then the labels would carry physical meaning not traceable to any recordable distinction, violating the admissibility constraint of Section 2. Under admissibility, the composition parameter t should mark the physical mixing ratio, not a label-dependent asymmetry. Condition (vi) is thus the operationalisation of label-neutrality for the composition parameter — it says that t marks the fractional weight of ψ_2 regardless of which state occupies which argument slot, which is a direct consequence of admissibility applied to the labelling of states in a composition. This provides a stronger physical grounding for (vi) than mere convention, though a full derivation from (A1) and (A2) alone is not given here.

Scope of Theorem 3.7. The theorem establishes that complete cancellation — $w(\psi_1 \star_t \psi_2, \lambda) = 0$ for some t and λ — is structurally necessary. It does not establish that such cancellation is generic, controlled, or basis-dependent in the specific way characteristic of quantum interference. Many non-classical theories (for example, theories with epistemic restrictions, or Popescu–Rohrlich box theories) exhibit suppression without the full phase-sensitive, basis-dependent pattern of quantum mechanics. The present result establishes the *minimum* structural condition: any non-classical theory must permit cancellation. The further constraint that this cancellation is governed by a linear amplitude framework with phase is supplied by Section 5; it is not derived from Section 3 alone.

Theorem 3.7 (Interference necessity). Any non-classical, regular composition rule admits cancellation between alternatives — and phase-sensitive cancellation in any linear embedding of the composition structure.

Proof. By Theorem 3.4, the composition rule is not a convex combination. By Proposition 3.6b, any composition rule satisfying support-monotonicity, invertibility, continuity, and closure must be affine — i.e. a convex combination. Since the rule is not a convex combination, it cannot satisfy support-monotonicity.

Therefore, there exist pre-factual states ψ_1, ψ_2 and a configuration λ such that:

$$w(\psi_1 \star_t \psi_2, \lambda) < \min(w(\psi_1, \lambda), w(\psi_2, \lambda))$$

for some t. By continuity (R1) and the intermediate value theorem, since $w \geq 0$ and the composed weight can fall below either component's weight, there exist parameter values at which $w(\psi_1 \star_t \psi_2, \lambda) = 0$ — that is, the contribution from λ is completely cancelled by the composition. This is the defining structural property of interference: composition can suppress, and in limiting cases eliminate, contributions supported by each component individually. In amplitude-based representations, this suppression is naturally encoded as phase-sensitive interference. □

Remark 3.8. The regularity conditions (R1)–(R3) are not derived from fact-production alone; they are additional structural requirements. (R1) and (R3) are motivated by the continuity and reversibility of pre-factual dynamics, both stated in Section 2. (R2) is the assumption that pre-factual composition is a genuine operation rather than merely sequential. These conditions should be understood as part of the specification of what "reversible pre-factual dynamics" means.

Remark 3.9 (Logical direction: interference before amplitudes). Theorem 3.7 establishes the necessity of *structural interference* — the ability of compositions to suppress contributions supported by individual alternatives — at the level of weights on configurations, prior to any introduction of amplitudes. The emergence of a phase representation is not assumed here and does not enter the proof. Phases and complex amplitudes arise later, in Section 5, from embedding this interference structure into a linear algebraic framework. The logical order is deliberate: interference is derived first as a structural property of the pre-factual sector; the amplitude formalism is the minimal algebraic language in which that structure can be expressed.

---

4. Necessity of Isotropy

4.1 Permutation Symmetry

Definition 4.1 (Equal distinguishability). Two configurations $\lambda_1, \lambda_2 \in \Lambda$ are *equally distinguishable* if $\delta(\lambda_1, \cdot) = \delta(\lambda_2, \cdot)$ as functions on Λ — that is, no admissible fact-producing operation distinguishes them by identity rather than by their relational properties.

Lemma 4.2 (Label irrelevance). If λ_1 and λ_2 are equally distinguishable, their labels carry no physical meaning.

Proof. A label is physically meaningful only if it is the content of some admissible record. If no admissible operation produces a fact distinguishing λ_1 from λ_2 by identity (as opposed to by relation to other configurations), then no fact can record label identity. By the admissibility constraint, label identity therefore carries no physical content. □

Theorem 4.3 (Permutation invariance). Observable predictions must be invariant under permutations of equally distinguishable configurations.

Proof. Let σ be a permutation of Λ with $\delta(\sigma(\lambda), \sigma(\lambda')) = \delta(\lambda, \lambda')$ for all $\lambda, \lambda' \in \Lambda$. If predictions changed under σ, the theory would assign different physical content to configurations that are relationally identical under σ — this would require a fact distinguishing them, which does not exist by Lemma 4.2. Hence predictions must be invariant. □

4.2 Extension to Continuous Isotropy

Theorem 4.3 yields discrete permutation symmetry. We now state the conditions under which this extends to continuous symmetry. These conditions are *additional assumptions*, not consequences of fact-production alone.

Proposition 4.4 (Conditions for continuous isotropy). Suppose in addition to the structure of Section 2:

- **(A3a)** *Metric homogeneity*: δ is a metric and the isometry group of (Λ, δ) acts transitively on Λ,
- **(A3b)** *Continuity of the pre-factual sector*: the action of the isometry group on $\mathcal{P}$ is continuous.

Then the admissibility argument of Theorem 4.3 extends to the full continuous isometry group, yielding **continuous isotropy**: observable predictions are invariant under all isometries of (Λ, δ).

Proof. Under (A3a), any two points in Λ are related by an isometry, hence equally distinguishable in the sense of Definition 4.1. Lemma 4.2 applies to every such pair. Under (A3b), the invariance argument of Theorem 4.3 applies continuously across the whole isometry group. □

Remark 4.5 (Honest scope). The result of this section is:

Admissibility alone ⇒ permutation invariance.
Admissibility + metric homogeneity + continuity ⇒ continuous isotropy.

We do not claim that fact-production alone forces continuous isotropy. Metric homogeneity (A3a) asserts that no configuration is intrinsically preferred over any other. This is physically motivated but not logically derived from the minimal premises. Deriving the geometry of Λ from more primitive conditions remains an open problem.

---

5. Division Algebra Structure and the Origin of Fields

We now derive the algebraic structure of the amplitude domain, replacing the earlier claim that it must be a field with the weaker and correct result that it must be a division algebra.

5.1 Amplitudes Without Born-Rule Presupposition

Definition 5.1 (Amplitude map). An amplitude map is a function $\alpha: \mathcal{P} \times \Lambda \to A$, where A is a set with a binary operation, assigning to each pre-factual state ψ and configuration λ an element $\alpha(\psi, \lambda) \in A$.

We require that A carry an *evaluation functional* $\nu: A \to [0, \infty)$ such that observable predictions (commitment probabilities) are determined by $\nu(\alpha(\psi, \lambda))$. For the subsequent arguments to be well-defined, ν is required to satisfy: (i) *continuity* — ν varies continuously with the amplitude, so that nearby amplitude maps yield nearby predictions. A second condition — *faithfulness* — will be imposed once the additive structure of A is established in Lemma 5.1a. Faithfulness requires that $\nu(a) = 0$ if and only if a is the additive identity of A (the null amplitude, which is the image under α of the null alternative $0 \in \mathcal{P}$ defined in (S1)); this ensures that distinct non-null amplitudes produce distinguishable predictions. Deferring faithfulness to after Lemma 5.1a avoids presupposing additive structure in A before it has been derived. We do *not* assume $\nu = |\cdot|^2$ or any specific functional form. The form of ν is to be derived; $|\cdot|^2$ arises only once we establish that A carries a norm-compatible involution (see Remark 5.5).

5.2 Compositional Requirements

We impose:

- **(C1) Associativity**: sequential composition of pre-factual evolutions composes associatively at the level of amplitudes.
- **(C2) Additivity under superposition**: the amplitude map is additive over superpositions of alternatives — once the superposition algebra of Lemma 5.1a is established, this extends to bilinearity in the usual algebraic sense once scalar multiplication is defined.
- **(C3) Invertibility**: reversible pre-factual evolution is invertible at the level of amplitudes — every non-zero amplitude has a compositional inverse.
- **(C4) Distributivity**: multiplication of amplitudes distributes over addition. This reflects the requirement that sequential evolution after a superposition is equivalent to the superposition of the sequentially evolved components.

Lemma 5.1a (Superposition algebra). Suppose the pre-factual sector supports:

- **(S1) Null alternative**: there exists a state $0 \in \mathcal{P}$ representing the absence of any alternative, satisfying $\psi \star 0 = \psi$ for all ψ,
- **(S2) Cancellation**: for any pre-factual state ψ, there exists a state $(-\psi)$ such that $\psi + (-\psi) = 0$ (the null alternative), which is required by the interference structure of Section 3 — cancellation to zero must be achievable,
- **(S3) Associative recombination**: superposition of alternatives is associative, i.e. $(\psi_1 + \psi_2) + \psi_3 = \psi_1 + (\psi_2 + \psi_3)$,
- **(S4) Exchange symmetry**: superposition is symmetric, i.e. $\psi_1 + \psi_2 = \psi_2 + \psi_1$.

Then the superposition operation induces an **abelian group structure** on A.

Proof. (S1) provides the identity element. (S2) provides additive inverses. (S3) provides associativity. (S4) provides commutativity. These are precisely the axioms of an abelian group. □

Remark 5.1b (Physical motivation and logical bridge from Section 3). (S3) and (S4) are natural requirements on recombination: the order and grouping of alternatives should not affect the result before commitment. (S1) reflects the logical possibility of no alternative being active.

(S2) requires careful treatment. Theorem 3.7 establishes that the composition of some alternatives can drive the weight at a given configuration to zero — that is, *some* cancellations are possible. This is not yet enough to derive (S2): knowing that certain pairs of alternatives can cancel does not immediately imply that *every* alternative admits a complete additive inverse. The bridge from Theorem 3.7 to (S2) runs as follows. Theorem 3.7 shows that for any ψ and any configuration λ with $w(\psi, \lambda) > 0$, there exist composition parameters and partner states such that the composed weight at λ reaches zero. If we require this cancellation to be achievable *universally* — for any ψ, by some element of the pre-factual sector — then the pre-factual sector must be closed under the operation of forming additive inverses. That closure is exactly (S2). The universality requirement is physically motivated by the fact-production framework itself: a theory in which some alternatives have no inverse would contain pre-factual states whose contributions could never be cancelled, making certain suppression patterns permanently unavailable. If fact-production requires that any alternative be in principle suppressible — that no pre-factual contribution is structurally immune from cancellation — then universality follows. The condition is therefore the *algebraic completion* of the universal cancellability implied by nontriviality and regularity: it takes the existential result of Theorem 3.7 and elevates it to a universal algebraic axiom. This is a genuine additional structural assumption — the regularisation of cancellation into a full group operation — and should be understood as such.

The specific physical motivation for requiring universality rather than mere possibility is as follows. If some alternative ψ lacked an additive inverse, the pre-factual sector would contain states that can never be fully cancelled by any partner — states that are, in this sense, irreducible. But irreducibility of a pre-factual state means it contributes inextinguishably to all future pre-factual compositions that include it: no physical process within the pre-factual sector could neutralise its contribution before commitment. This would constitute a structurally privileged alternative — one that carries physically distinguishable weight regardless of the composition history — in tension with the admissibility constraint (A2), which forbids distinctions that no fact can record. A pre-factual state that cannot be cancelled is one whose presence is, in principle, detectable from the pre-factual sector alone, without commitment to a fact. Requiring that every alternative admit an additive inverse is therefore the condition that no alternative is irreducibly distinguishable within the pre-factual sector itself.

Theorem 5.2 (Division algebra structure). If the amplitude domain A satisfies Lemma 5.1a together with (C1)–(C4), is closed under addition and multiplication, and is **finite-dimensional over** $\mathbb{R}$, then A carries the structure of a **division algebra** over a subfield of the reals, and the only such algebras are $\mathbb{R}$, $\mathbb{C}$, and $\mathbb{H}$.

Proof. By Lemma 5.1a, (A, +) is an abelian group. (C1) and closure under multiplication make (A, ·) a monoid with identity (the amplitude corresponding to the identity evolution). (C3) makes every non-zero element invertible under multiplication — since reversible dynamics must be invertible at the amplitude level, and a non-invertible non-zero amplitude would represent a physically irreversible pre-factual evolution. (C4) ensures left and right distributivity of multiplication over addition.

A structure with an abelian additive group, an associative and invertible multiplication, and two-sided distributivity is a *division ring* (skew field). If the centre of A — the sub-ring of elements commuting with all others under multiplication — contains a copy of $\mathbb{R}$, then A is a division algebra over $\mathbb{R}$.

By the Frobenius theorem — which applies precisely to *finite-dimensional* division algebras over $\mathbb{R}$ — the only such algebras are $\mathbb{R}$, $\mathbb{C}$, and $\mathbb{H}$ (the quaternions). The finite-dimensionality assumption is physically motivated: any pre-factual sector with a finitely distinguishable configuration space Λ will induce a finite-dimensional amplitude domain; an infinite-dimensional A would require infinitely many mutually distinguishable configurations, which is a stronger structural assumption not needed here. □

Remark 5.3 (Division algebra vs. field). A field requires commutativity of multiplication. The derivation above does not yield commutativity automatically — sequential composition of physical processes is generically non-commutative. We therefore obtain a *division algebra*, not a field, without additional assumptions. This matters: $\mathbb{H}$ is a division algebra but not a field. Earlier versions of this paper incorrectly referred to all three candidates as "fields"; the present treatment corrects this.

Proposition 5.4 (Commutativity condition). The amplitude domain A is a commutative field if and only if, in addition to (C1)–(C4):

- **(C5) Commutativity**: for all a, b ∈ A, ab = ba.

Whether (C5) holds is a question about the structure of the pre-factual dynamics. It is *not* derivable from fact-production alone. We note that $\mathbb{R}$ and $\mathbb{C}$ satisfy (C5); $\mathbb{H}$ does not. The selection among these three is addressed in Section 7.

Remark 5.5 (Born rule). Once A is identified as $\mathbb{C}$, the evaluation functional ν can be identified with the modulus squared $|\cdot|^2$ via the inner product structure of complex Hilbert space. This is not assumed here but follows from the identification A = $\mathbb{C}$ and the geometry of the resulting state space.

---

6. Necessity of Representational Invariance and Galois Invariance

Definition 6.1 (Representational redundancy). A transformation φ: A → A of the amplitude domain is *redundant* if it preserves the evaluation functional ν and all observable relations — that is, ν(φ(a)) = ν(a) for all a ∈ A, and the transformed amplitude map α' = φ ∘ α produces identical predictions to α.

Lemma 6.2 (Elimination of redundancy). A physical theory must identify all redundant descriptions.

Proof. If two descriptions yield identical predictions for all admissible observations, no fact can distinguish them. By admissibility, they carry the same physical content. A theory treating them as distinct introduces non-recordable structure, violating admissibility. □

Theorem 6.3 (Representational invariance). Observable predictions must be invariant under all redundant transformations of the amplitude domain.

Proof. Immediate from Lemma 6.2 and Definition 6.1. □

Corollary 6.4 (Galois invariance — conditional on commutativity). If the amplitude domain A is a commutative field (i.e. (C5) holds in addition to (C1)–(C4)), representational invariance reduces to invariance under field automorphisms — that is, under the Galois group of A over the subfield fixed by ν.

Proof. A field automorphism of A preserving ν satisfies Definition 6.1 and is therefore redundant. Theorem 6.3 requires predictions to be invariant under all such automorphisms, which form the Galois group. □

Corollary 6.5 (Division algebra case). If A is a non-commutative division algebra (e.g. ℍ), representational invariance applies to all ν-preserving automorphisms of A as a division algebra. This is a weaker condition than Galois invariance in the field-theoretic sense.

Remark 6.6 (Logical structure). The derivation proceeds: division algebra structure (Theorem 5.2) → identification of ν-preserving automorphisms → representational invariance (Theorem 6.3) → Galois invariance if and only if (C5) holds (Corollary 6.4). Galois invariance is not an independent postulate; it is a consequence of representational invariance restricted to the commutative case.

---

7. Selection of Complex Amplitudes

We now address the selection among ℝ, ℂ, and ℍ. We are explicit about which arguments are self-contained here and which depend on prior work.

7.1 Ruling Out ℝ: Failure of Continuous Isotropy for Composite Systems

Claim 7.1. Under the transitivity result established in [Taylor, VERSF Distinguishability Geometry, Theorem X], real quantum mechanics fails the continuous isotropy condition (Proposition 4.4) for systems of dimension $n \geq 3$.

Argument. In real quantum mechanics, the state space of a single system of dimension n is the real projective space ℝP^(n−1), and the isometry group under the Fubini–Study analogue metric is the projective orthogonal group PO(n).

For $n \geq 3$, PO(n) does not act transitively on *pairs* of orthogonal states: the orbit of a pair (λ_1, λ_2) with $\delta(\lambda_1, \lambda_2) = d$ under PO(n) does not exhaust all pairs at distinguishability d when the ambient dimension is $n \geq 3$ and the pair carries additional relational structure (such as being jointly orthogonal to a third state). Consequently, there exist configurations that are equally distinguishable pairwise but not related by any isometry — the isometry group is not transitive on the full set of equally distinguishable pairs in the presence of a third system.

Dependence on prior work: The precise transitivity analysis for the Fubini–Study geometry and its real analogue — specifically, the failure of PO(n) to act transitively on triples of mutually orthogonal states for $n \geq 3$ — is established in [Taylor, VERSF Distinguishability Geometry, Theorem X]. The present paper establishes the conditional: transitivity failure implies violation of Proposition 4.4, and hence of continuous isotropy. The transitivity result itself is imported and not reproved here.

Consequence. Under the admissibility argument, configurations that are equally distinguishable but not related by an isometry could in principle receive different physical treatment. Real quantum mechanics introduces such structure at $n \geq 3$, violating the isotropy requirement.

7.2 Ruling Out ℍ: Failure of Representational Invariance for Composite Systems

Claim 7.2. Quaternionic quantum mechanics fails representational invariance (Theorem 6.3) when applied to composite systems.

Argument. In quaternionic quantum mechanics, the tensor product of two quaternionic Hilbert spaces $\mathscr{H}_1 \otimes\_\mathbb{H}\ \mathscr{H}_2$ does not admit a canonical definition compatible with local operations. There exist multiple inequivalent ways to define bipartite states that agree on all single-system marginals. This is a consequence of the non-commutativity of ℍ: the module structure of $\mathscr{H}_1 \otimes\_\mathbb{H}\ \mathscr{H}_2$ depends on a choice of how to extend the right-module structure of $\mathscr{H}_1$ across the tensor product, and different choices yield inequivalent bipartite observable algebras.

More precisely: let α and α' be two bipartite amplitude maps differing by a transformation within the non-commutative endomorphism algebra of the quaternionic module — specifically, one arising from the non-central part of that algebra. Then α and α' agree on all single-system observables (since those are insensitive to the bipartite module structure) but differ on genuinely bipartite observables. The transformation between them satisfies $\nu(\varphi(a)) = \nu(a)$ on single systems but alters bipartite predictions. This is precisely the type of non-redundant representational ambiguity ruled out by Theorem 6.3.

Dependence on prior work: The detailed analysis of quaternionic tensor products and the non-uniqueness of bipartite structure is developed in [Taylor, VERSF Distinguishability Geometry, Theorem Y], which establishes the specific module-inequivalence result. The present paper establishes that such inequivalence constitutes a violation of representational invariance; the quaternionic analysis is imported.

Remark on C5 and the exclusion of ℍ. The exclusion of ℍ via this argument does not require condition (C5) — it does not depend on the amplitude domain being a commutative field. The relevant principle is Theorem 6.3 (representational invariance) and its Corollary 6.5, which applies to non-commutative division algebras including ℍ. The argument is that the specific non-central automorphisms of the quaternionic module algebra are ν-preserving on single systems but not globally redundant, because they alter bipartite predictions. This violates Theorem 6.3 directly, without needing to invoke Galois invariance (Corollary 6.4) or commutativity. The exclusion of ℍ thus stands independently of the unresolved status of (C5).

Consequence. ℍ introduces representational non-redundancy at the level of composite systems, violating Lemma 6.2.

7.3 Complex Amplitudes Satisfy All Constraints

ℂ satisfies all derived constraints:

- It is a commutative field, satisfying (C5) and yielding Galois invariance via Corollary 6.4.
- The projective space ℂP^(n−1) with the Fubini–Study metric admits U(n) as its isometry group, which acts transitively on all pairs of states at equal distinguishability for all n. This satisfies metric homogeneity (A3a) for all system dimensions, unlike ℝ.
- The tensor product of complex Hilbert spaces is canonically defined and compatible with local operations, introducing no non-redundant bipartite structure. This satisfies representational invariance for composite systems, unlike ℍ.

The selection of ℂ is therefore a consequence of simultaneously satisfying isotropy at all dimensions and representational invariance for composite systems. The full proof that ℂ is the *unique* solution — ruling out other division algebras or exotic structures not covered by the Frobenius theorem — is completed in [Taylor, VERSF Distinguishability Geometry, Theorem Z], which establishes uniqueness via distinguishability geometry.

---

8. The Bridge Theorem

Theorem 8.1 (Bridge Theorem). Let a physical theory satisfy:

- **(A1)** A nontrivial reversible pre-commitment sector,
- **(A2)** Admissibility of physical distinctions,
- **(A3)** Metric homogeneity and continuity of the isometry action on $\mathcal{P}$,
- **(A4)** Associative, bilinear, invertible, distributive composition of alternatives,
- **(A5)** Locality of composition: weight updates at a configuration depend only on prior weight at that configuration and the composition parameter.

Then:

Result	Status
(1) Non-classical composition required	**Established here** — Theorem 3.4
(2) Structural cancellation required; phase-sensitive in linear embedding	**Established here** — Theorem 3.7, under regularity (R1)–(R3) and (A5)
(3) Permutation invariance	**Established here** — Theorem 4.3, from (A1)–(A2)
(4) Continuous isotropy	**Established here** — Proposition 4.4, from (A2)–(A3)

Result	Status
(5) Division algebra structure	**Established here** — Theorem 5.2, from (A4)
(6) Representational / Galois invariance	**Established here** — Theorem 6.3 and Corollary 6.4
(7) Exclusion of $\mathbb{R}$ and $\mathbb{H}$	**Conditional on prior work** — Section 7 identifies which constraint fails and establishes the conditional; transitivity and module-inequivalence results imported from prior work
(8) Unique determination of complex Hilbert space	**Imported** — prior work on distinguishability geometry

Remark 8.2 (Non-circularity). The derivation does not presuppose amplitudes, Hilbert space, or quantum structure. Amplitudes enter as a representation of pre-factual states (Definition 5.1), required by nontriviality. The evaluation functional ν is introduced without specifying its form. Division algebra structure is derived from compositional requirements (C1)–(C4). Galois invariance follows from representational redundancy, which is itself a consequence of admissibility. The residual dependence on prior work is in (7)–(8): the selection of $\mathbb{C}$ over alternatives, and the final uniqueness result.

Remark 8.3 (Status of (A3), (A4), and (A5)). These three assumptions go beyond the minimal fact-production premise:

- **(A3)** — metric homogeneity — asserts that the configuration space has no intrinsically preferred points. It is motivated by admissibility but not derived from it. Whether the geometry of Λ can be derived from more primitive conditions remains an open question.
- **(A4)** — the compositional requirements (C1)–(C4) — are motivated by the structure of reversible dynamics and superposition but require independent justification. In particular, distributivity (C4) and the additive group structure of superposition are assumed rather than derived from the bare fact-production premise.
- **(A5)** — locality of composition — asserts that weight updates at a configuration are independent of the detailed distribution over other configurations. It is physically natural and implicit in the semigroup law (v) under uniform application, but is stated as an explicit assumption rather than derived.

---

9. Responses to Anticipated Objections

The following addresses objections likely to arise in peer review, organised by target section.

O1. "The fact-production premise is too vague to bear formal weight." (Sections 2–3)

The premise — that a physical universe must produce stable, irreversible facts — is intentionally minimal. It is not a full theory of measurement or ontology. Its formal content is captured entirely by Definition 2.3 (nontriviality): the pre-factual sector must be operationally distinguishable from a classical mixture. Everything downstream follows from that formal condition plus the explicitly stated additional assumptions (A1)–(A5). A reader who rejects the premise is invited to specify which physical universe they have in mind that produces no irreversible records whatsoever; the claim is that any framework admitting such records will fall under the present conditions.

O2. "(A5) is a classicality assumption smuggled into the interference derivation." (Lemma 3.6b')

This objection is addressed directly in the compatibility remark following the definition of (A5). The key distinction is between two levels of description. (A5) governs the marginal update of weight at a configuration under composition with a *definite* alternative — the probability level. The explicit QM verification (Lemma 3.6b' remark) shows that quantum mechanics itself satisfies (A5) at this level: $\mathrm{Tr}(\rho_t \Pi_{\{\lambda_1\}}) = (1-t)\cdot p(\lambda_1)$ when $\rho_t = (1-t)\rho + t|\lambda_2\rangle\langle\lambda_2|$. Quantum interference lives in the off-diagonal terms of ρ and manifests at the amplitude level, which (A5) does not constrain. Far from assuming classicality, (A5) is a condition that quantum mechanics satisfies while still being fully non-classical at the amplitude level.

O3. "Proposition 3.6b only shows cancellation is possible, not that the composition is structurally non-classical in the quantum-mechanical sense." (Theorem 3.7)

Correct, and the paper makes this explicit in Remark 3.6c. Theorem 3.7 establishes the *minimum* structural condition: any non-classical theory must permit cancellation. PR-box theories, epistemic-restriction theories, and other non-classical frameworks also satisfy this minimum. The quantum-specific structure — phase-sensitive, basis-dependent cancellation governed by a linear amplitude framework — is not claimed here; it is derived in Section 5 from the division algebra argument. The table in Section 8 reflects this explicitly: row (2) reads "structural cancellation required; phase-sensitive in linear embedding."

O4. "Conditions (v) and (vi) of Proposition 3.6b are load-bearing but not derived." (Proposition 3.6b)

Acknowledged in Remark 3.6c, which now provides an admissibility-based motivation for each. Condition (v) (the semigroup/no-memory property) reflects that the pre-factual sector has no hidden memory of intermediate composition steps — only the compound degree matters, not the path. Condition (vi) (exchange symmetry) is motivated by admissibility directly: if labelling ψ_1 "first" versus ψ_2 "first" produced different observable predictions, those labels would carry physical meaning not traceable to any recordable distinction, violating (A2). Both conditions are structural assumptions that go beyond the bare fact-production premise; the paper is explicit that they are not derived and lists their status in Remark 8.3.

O5. "The jump from Theorem 3.7's cancellation to (S2)'s universal additive inverses is not proved." (Lemma 5.1a)

This gap is explicitly acknowledged and addressed in Remark 5.1b. Theorem 3.7 shows cancellation is possible for some alternatives; (S2) requires it universally. The bridge argument is as follows. A pre-factual state ψ lacking an additive inverse would carry irreducible weight in every composition that includes it — no partner state could neutralise its contribution before commitment. This irreducibility means that ψ's presence affects commitment outcomes differently from a world in which ψ were cancellable: the pattern of facts produced by the universe would systematically differ depending on whether ψ can be cancelled, and this difference would be recordable in the committed outcomes. The alternative ψ therefore carries physically distinguishable weight that manifests at the level of facts, even though its inextinguishability is a property of the pre-factual sector. This is the sense in which an uncancellable alternative violates the admissibility constraint (A2): the distinction is ultimately recordable at commitment, not merely detectable in the reversible pre-factual dynamics. (S2) is the structural condition that prevents any alternative from occupying this privileged, inextinguishable role. It is a genuine additional assumption — the algebraic regularisation of cancellation into a full group operation — and is labelled as such in Remark 5.1b.

O6. "The exclusion of ℍ depends on commutativity (C5), which is not derived." (Section 7.2)

It does not. The exclusion of ℍ runs through Theorem 6.3 (representational invariance) and Corollary 6.5, which applies to non-commutative division algebras without invoking (C5) or Galois invariance. The argument is that the non-central automorphisms of the quaternionic module algebra are ν-preserving on single systems but alter bipartite predictions — this is a direct violation of Theorem 6.3, not mediated by the field/non-field distinction. This is clarified in the remark on (C5) in Section 7.2. The unresolved status of (C5) does not compromise the ℍ exclusion.

O7. "Continuous isotropy is not derived from fact-production — you need metric homogeneity as an extra assumption." (Proposition 4.4)

Correct, and the paper says so explicitly in Remark 4.5. Admissibility alone yields permutation invariance (Theorem 4.3). Continuous isotropy requires additionally that (Λ, δ) be metrically homogeneous (A3). Metric homogeneity asserts that no configuration is intrinsically preferred — it is the geometric expression of the same admissibility principle, applied to the continuous structure of Λ rather than to discrete relabellings. Whether metric homogeneity can be derived from more primitive conditions is flagged as open question 2 in the Discussion.

O8. "The paper relies on placeholder citations for key results in Section 7." (Section 7)

Acknowledged in the Discussion's note on citations. The placeholders [Taylor, VERSF Distinguishability Geometry, Theorems X, Y, Z] will be replaced with full references on publication of the companion paper. The present paper's contribution is clearly demarcated: it identifies *which* constraint each candidate fails (isotropy for $\mathbb{R}$, representational invariance for $\mathbb{H}$) and establishes the conditional (transitivity failure implies isotropy violation; module non-uniqueness implies representational invariance violation). The positivity of the identification is established here; the underlying geometric results are imported. This division of labour is reflected in Table 8.1.

O9. "The definition of ν in Section 5.1 is too permissive to constrain the algebra." (Definition 5.1)

Two explicit conditions are now placed on ν: continuity (imposed immediately) and faithfulness (imposed after Lemma 5.1a establishes the additive structure, to avoid circularity). Continuity ensures predictions vary smoothly with amplitudes; faithfulness ensures null amplitudes are the only zero-prediction amplitudes. These are the minimum requirements needed for the division algebra and representational invariance arguments. The paper does not assume $\nu = |\cdot|^2$, which is instead derived as a consequence once A is identified as $\mathbb{C}$ via the full programme.

O10. "Theorem 5.2 proves only a division algebra, not complex structure specifically. The result is too weak." (Theorem 5.2)

This is correct and intentional. Theorem 5.2's purpose is to narrow the candidate amplitude domains to $\{\mathbb{R}, \mathbb{C}, \mathbb{H}\}$ via the Frobenius theorem — a substantial narrowing from the unconstrained set of all possible algebraic structures. The further selection among these three is the task of Section 7, which uses isotropy (to exclude $\mathbb{R}$) and representational invariance (to exclude $\mathbb{H}$). The full derivation of complex structure requires both Theorem 5.2 and Section 7; neither is claimed to suffice alone. The paper is structured to make this two-stage argument explicit.

---

10. Discussion

The results clarify the foundational status of quantum mechanics in three ways.

A note on citations. Several results in Section 7 are imported from prior VERSF work, currently cited as [Taylor, VERSF Distinguishability Geometry, Theorems X, Y, Z]. These placeholders will be replaced with full bibliographic references upon publication of that companion paper. Readers should treat the imported results as clearly demarcated dependencies rather than as established corollaries of the present derivation.

Reduction of independent assumptions. The structural principles used in previous work — interference, isotropy, Galois invariance — are here shown to follow from four conditions (A1)–(A4), together with the locality assumption (A5). These conditions are more primitive and more transparently motivated than the structural principles themselves. The logical chain from fact-production to quantum structure is now explicit, with every dependency labelled. (A5) is identified in the open questions as the most important assumption not yet derived from more primitive conditions; resolving its status is the central outstanding task for the programme.

Honest scope. We have been careful to distinguish what is derived here from what is imported. The present paper establishes the *necessity* of non-classical, phase-sensitive composition, permutation symmetry, division algebra structure, and representational invariance. It identifies which constraints exclude $\mathbb{R}$ and $\mathbb{H}$, but relies on prior distinguishability geometry for the full uniqueness result.

Interpretation of quantum structure. Quantum mechanics is not an arbitrary framework. It is the minimal structure compatible with a universe in which unresolved possibilities can become irreversible facts, without introducing unobservable distinctions at the level of configurations, representations, or composite systems.

Remaining open questions. The derivation leaves the following open:

1. *Locality of composition (A5)*: this is the most important unresolved assumption in the paper. (A5) asserts that marginal weight updates under composition with a definite alternative depend only on prior weight and the composition parameter. While physically natural and consistent with quantum mechanics at the probability level, a derivation of (A5) from (A1)–(A4) alone has not been given. Whether (A5) follows from the other assumptions under additional regularity conditions — or whether it constitutes an independent postulate representing a genuine constraint on the class of physically admissible composition rules — is the central open question for strengthening the present derivation.
2. *Geometry of* Λ: metric homogeneity is assumed, not derived. A deeper account would derive the manifold structure of the configuration space from fact-production constraints.
3. *Dynamics*: we derive the kinematic structure but not the dynamical law. Unitarity and the Schrödinger equation remain to be derived.

4. *Commutativity of multiplication*: condition (C5) is not derived from (A1)–(A5). Its justification remains incomplete, though it does not affect the exclusion of ℍ (see Section 7.2).
5. *Dimensionality*: the present framework does not fix the Hilbert space dimension.

---

11. Conclusion

We have shown that any physical theory capable of producing stable, irreversible facts from unresolved alternatives must satisfy:

- **Non-classical, phase-sensitive composition** (from nontriviality and regularity),
- **Permutation symmetry and, conditionally, continuous isotropy** (from admissibility and metric homogeneity),
- **Division algebra structure for amplitudes** (from associativity, bilinearity, distributivity, and invertibility),
- **Representational invariance and Galois invariance** (from elimination of unobservable structure).

These are structural requirements imposed by the existence of irreversible facts, not independent postulates. Using the structural exclusions identified here — ℝ fails continuous isotropy at $n \geq 3$, ℍ fails representational invariance for composite systems — together with the transitivity and uniqueness results established in prior work, ℂ remains the unique surviving candidate.

The core mathematical structure of quantum mechanics is forced by the requirement that reality produce facts without unobservable distinctions — subject to the explicitly stated assumptions and imported results identified throughout.

Appendix: Logical Dependency Map

```
Fact-production requirement
              │
              ├──→ Nontriviality (A1)
              │          │
              │          └──→ Non-classical composition (Theorem 3.4)
              │                          │
              │                          └──[+ regularity (R1)-(R3)]
              │                                         │
              │                              Lemma 3.6b': local attenuation
              │                              [+ (A5) locality of composition]
              │                              w factorises multiplicatively at each
λ
              │                                         │
              │                              Prop 3.6b: affine reduction
```

```
           |                             [+ semigroup law (v) + exchange
symmetry (vi)]
           |                             Cauchy equation → p(t)=(1-t)^c → c=1
           |                                      |
           |                                      └──→ Structural interference
(Theorem 3.7)
           |                                             [phase-sensitive in
amplitude representations]
           |
           ├──→ Admissibility (A2)
           |          |
           |          ├──→ Permutation invariance (Theorem 4.3)
           |          |          |
           |          |          └──[+ metric homogeneity (A3)]
           |          |                        |
           |          |                        └──→ Continuous isotropy
(Proposition 4.4)
           |          |
           |          └──→ Representational invariance (Theorem 6.3)
           |                          |
           |                          └──[+ commutativity (C5)]
           |                                      |
           |                                      └──→ Galois invariance
(Corollary 6.4)
           |
           └──→ Compositional structure (A4)+(A5)
                      |
                      ├──→ Lemma 5.1a: abelian group on A
                      |         [(S1)-(S4): null, algebraic completion
                      |          of cancellation, associativity, exchange]
                      |          |
                      └──────────┴──→ Division algebra (Theorem 5.2)
                                            |
                                            └──[+ (C5)]──→ Field structure
(Proposition 5.4)

[Interference + isotropy + field structure + Galois invariance]
           |
           ├──→ Exclude ℝ (fails isotropy, n ≥ 3)
           |          [Section 7.1 + Taylor, VERSF Dist. Geom., Thm X]
           ├──→ Exclude ℍ (fails representational invariance)
           |          [Section 7.2 + Taylor, VERSF Dist. Geom., Thm Y]
           └──→ Select ℂ
                      [Section 7.3 + Taylor, VERSF Dist. Geom., Thm Z:
uniqueness]
                      |
                      └──→ Complex Hilbert space [prior work]
```
