Observer-First Physics

A Derivation of Quantum Mechanical Structure
from the Conditions of Observation

J.B. Thompson

Abstract

We derive the full mathematical structure of finite-dimensional quantum mechanics—complex Hilbert space, tensor-product composition, the Born rule, and unitary dynamics—from five observer constraints and a single invariance principle. The observer constraints are: physical situatedness (observers are finite systems inside the world they observe), finite causally ordered memory, bounded information storage, interaction-only epistemology, and the identification of physically indistinguishable histories. The invariance principle, Ground Axiom A0, states that physical reality consists solely of what is invariant under physically inert transformations—pure redescription. The central technical result is the redecomposition argument: because any orthogonal rotation of the description basis of a system is pure redescription, the distinguishability measure must be O(n)-invariant, uniquely forcing the Euclidean norm, the inner product, and Hilbert space structure. Group structure on the state space is derived via a cancellativity lemma and the Grothendieck construction, making the structure theorem for locally compact abelian groups applicable. Translation-invariance of the operational metric is derived from causal independence, bridging the operational metric to the quadratic norm. The phase group U(1) is derived from the requirement that free dynamics of isolated systems preserve probability assignments while changing physical states—a consequence of interaction-only epistemology, not an assumption about quantum structure. The qubit Born-rule gap is closed internally: the observer's presence guarantees every physical system appears in a joint space of dimension ≥ 4. No alternative structure survives all constraints simultaneously. The derivation is finite-dimensional throughout; the

infinite-dimensional limit is identified as the natural open problem.

Contents

Declaration

Every reader of this paper is a finite physical system inside the world it describes.

This is not a philosophical preliminary. It is the operative fact from which everything that follows is derived. You are made of matter. You arrived here through physical processes. Every piece of information you have ever possessed reached you through physical interaction. You cannot step outside the world to examine it from a neutral vantage point. Neither can we.

This paper asks: given that every observer is constitutively inside the world—not visiting it, not modeling it from outside, but *being* a physical part of it—what must the structure of physics look like?

The answer is not that quantum mechanics is one possible structure that happens to fit the data. The answer is that quantum mechanics is the *unique* structure consistent with observation being what it actually is: a finite, local, interaction-based, physically situated process. The Hilbert space, the Born rule, unitary dynamics, entanglement—none of these are features discovered by observers looking at the world. They are features constituted by what observation necessarily is when the observer is inside the world.

The derivation that follows is formal and gap-named. Its claim is that the apparent mystery of quantum mechanics—why complex amplitudes, why the Born rule, why the tensor product—dissolves entirely once you stop treating the observer as external to the physics and recognize that physical existence, for any system accessible to an observer inside the world, just *is* the invariant structure of its interactions.

We exist. We observe. This is existence, and you are already in it. What follows is the proof that this recognition forces, uniquely and completely, the structure of quantum mechanics.

This Paper and the Reconstruction Literature

This paper is not a reconstruction program in the sense of Hardy [3], Chiribella–D'Ariano–Perinotti [1], or the generalized probabilistic theories (GPT) framework. Those programs ask: what axioms, imposed on an abstract state-space formalism, uniquely select quantum mechanics? They start with a formalism and constrain it from outside.

This paper asks a different question: what structure is *constituted* by the physical situation of a finite observer inside the world it observes? It starts with the observer's actual physical situation and derives the formalism. Ground Axiom A0 is not a constraint imposed on a pre-existing formalism — it is a recognition about what physical existence means for systems that learn only through interaction. The derivation does not assume quantum structure and derive it from axioms imposed above; it shows that quantum structure is what the physical situation of observation necessarily *is*.

This distinction matters for how to read the argument. The question to ask of each step is not "is this axiom independently motivated from outside?" but "does this follow from what it means to be a finite physical system inside the world?" A0 is not a universal solvent for unexplained structure; it is the constitutive recognition that every reader has already made by virtue of being a finite physical system reading this. The derivation makes that recognition explicit and follows it to its unique mathematical conclusion.

A specific objection must be addressed directly: "you defined a symmetry class whose fixed points are quantum mechanics, then concluded quantum mechanics is unique." This objection assumes A0 is selecting the symmetry group strategically. It is not. A0 makes a single constitutive claim: physically inert transformations have no physical content. The symmetry group is not specified by A0 — it is identified by the subsequent steps, which ask what transformations are physically inert given the constraints of finite embedded observers. That the answer turns out to be the symmetry group of quantum mechanics is the result, not the assumption. The charge of circularity would require showing that A0 covertly encodes quantum-specific symmetry structure; no such encoding is present. A0 applies equally to any physical theory; it is the observer constraints (Steps 1–5) that select quantum mechanics as the unique fixed point.

Ground Axiom A0: Physics is Invariant Structure

The gauntlet. To reject A0 is to assert that there exist physical facts that make no difference to any outcome of any physical process, ever, under any protocol, for any observer, anywhere. Facts that leave no trace, produce no effect, enter no prediction, change no experiment, and yet are real. The burden of proof is not on the person who says such things do not exist. It is on the person who says they do. This is the challenge A0 throws down, and it has not been met in any known case.

The physical cogito. The argument begins not with a postulate but with a fact about the reader.

You are reading this as a finite physical system. That is not a philosophical framing — it is a description of what is happening right now. You are made of matter. Your reading is a physical interaction between you and this text. The uncertainty you may feel about any claim here is a state change in a finite material system. You cannot step outside this situation to evaluate it from neutral ground. Neither can we. Neither can anyone.

This is the physical version of the cogito. Descartes showed that the act of doubting one's existence is itself the proof of it — doubt is a form of thinking, and thinking is existence. The physical version is harder and stronger: you cannot doubt that existence is constituted by interaction without performing an interaction. The act of doubt is a physical process. The instrument of skepticism is made of the same stuff as the thing it questions. You are already inside A0. You were inside it before you read the first word.

Extending to the collective: we exist, we observe, and what we find is that physical existence just is the invariant structure that survives every possible redescription by interacting observers. Not because we are blind to anything beyond it. Because physical existence for a system inside the world *is* constituted by interaction. The physical world is exhausted by the invariants of interaction by constitution, not by ignorance. The distinction is between blindness and constitution, and it is not a subtle one.

What A0 is not. A0 is not verificationism. Verificationism is epistemic: we cannot detect X, therefore X does not exist. A0 is ontological: physical existence for embedded systems just is invariant interaction structure, so non-invariant "structure" is not a limitation of our access but an absence of physical content. Verificationism could be wrong if there were physical facts we merely happen to be unable to detect. A0 cannot be wrong in that way, because it does not depend on the limits of current detection — it depends on the constitution of physical existence for systems that are themselves physical.

The historical record. Every time this gauntlet has been thrown in the history of physics, physics has advanced by picking it up. Leibniz challenged Newton on absolute space: what experiment distinguishes absolute position from relative position? None. The distinction had no physical content. It was discarded and mechanics was clarified. Mach pressed the same challenge on absolute rotation. Einstein followed it to special relativity by asking what experiment could ever detect absolute simultaneity. None. Simultaneity was frame-dependent and the universe did not care about the convention. The invariant — the spacetime interval — was the physics. A0 is the same move brought to the foundations of observation itself. The invariant structure of physical interaction is the physics. Everything else is coordinate choice.

Axiom 1 (A0: Physics is invariant structure)**.** Physical reality consists solely of what is invariant under all physically inert transformations. A transformation is physically inert if and only if it changes no outcome of any physical process for any observer — it is pure redescription. Any purported physical fact that makes no difference to any outcome of any physical process, under any protocol, for any observer, is not a physical fact.

A0 is not derived from the observer steps in the narrow logical sense. It is the recognition that makes those steps self-consistent and non-circular. Every observer is a finite physical system certified by interaction. That certification extends: the world those observers inhabit is the invariant structure of their interactions, because there is nothing else their existence could be constituted by. A0 is what falls out when you take the physical situation of the observer seriously all the way down.

Scope. The framework describes physics accessible to observers satisfying the steps below. Physical systems or regions that never interact with any such observer are outside the domain of the derivation. This is a statement about the structure of physics as it necessarily presents itself to finite embedded observers, not a claim about the totality of the universe.

Consequence. The state space is the orbit space of all descriptions under physically inert transformations. Indistinguishability under admissible protocols is the operational instrument for certifying that a transformation is inert. The foundation is the invariance; indistinguishability is how we certify it.

The Redecomposition Argument

Splitting a physical system into independent parts and recombining them is redescription. It changes no physical process, no outcome, no interaction. By A0 it changes nothing physically real—including the total distinguishability of the system from other systems.

The combination rule for distinguishability must therefore be invariant under all admissible splittings, including splits at arbitrary angles. The key physical claim here: *any choice of orthonormal basis is an equally valid description of the same physical state.* Basis choice carries no physical content by A0—it is pure convention. Therefore the norm, which encodes physical distinguishability, must be invariant under all orthonormal basis changes. These are exactly the orthogonal transformations $\mathrm{O}(n)$.

Lemma 1 (Redecomposition forces the Euclidean norm)**.** *The unique norm on $\mathbb{R}^n$ invariant under all orthogonal transformations is the Euclidean norm. Any p-norm with $p \neq 2$ fails this invariance.*

Proof. For the 45° rotation R_{45} in $\mathbb{R}^2$, direct calculation:

$$\|R_{45}(1,0)\|_p = \left(\left(\frac{1}{\sqrt{2}}\right)^p + \left(\frac{1}{\sqrt{2}}\right)^p\right)^{1/p} = 2^{1/p-1/2}$$

which equals 1 only when $p = 2$. For $p = 3$: $2^{-1/6} \approx 0.891$, a 10.9% deviation. For $p = 1$: $2^{1/2} \approx 1.414$. A composite protocol executing this split detects a norm change for any $p \neq 2$, making the split physically detectable and violating A0. The unique $\mathrm{O}(n)$-invariant norm on $\mathbb{R}^n$ is the Euclidean norm [10]. □

This lemma will be applied in Step 10, once the vector space structure of the state space has been derived. It is stated here for motivation; its application is formal at Step 10.

For independent systems x, y in orthogonal subspaces, the forced combination rule is:

$$D(x \oplus y)^2 = D(x)^2 + D(y)^2$$

The Derivation Chain

Step 1: Observers are physical things

Every observer is a physical system inside the world it observes—made of matter, consuming energy, occupying space, arrived here through physical processes. All information arrives through physical interaction. There is no external channel.

$$O \in S, \quad \text{finite mass}, \quad E < \infty$$

Step 2: Memory is finite and causally ordered

$$H = (o_1, o_2, \ldots, o_n), \quad o_i \in O, \quad n \leq N < \infty$$

Strict causal ordering from the physical process of recording. Finite bound from finite storage.

Primitive relation. Two interaction events are *causally independent* if neither precedes the other in this causal ordering. Defined by causal structure alone, prior to any equivalence relation on histories.

Step 3: Information storage is bounded

$$2 \leq \dim(S) \leq k < \infty$$

Upper bound: Bekenstein bound $S \leq (2\pi k_B RE)/(\hbar c)$ [4]. Lower bound: the observer must store at least one bit to record any outcome at all. Both follow from Step 1 via established physics.

Step 4: Only interactions produce physical knowledge

$$S_O = \{\text{states reachable by finite interaction sequences from } O\}$$

$$\text{Prediction} = F(H)$$

Step 5: The physical state space is the invariant quotient

By A0, the physical state space is the orbit space of histories under physically inert transformations. The operational criterion: no admissible protocol detects the transformation.

Operational metric. Over all admissible protocols π constructible under Steps 1–4:

$$d(H_1, H_2) = \sup_{\pi \text{ admissible}} D_{\mathrm{TV}}\big(X_{H_1}(\pi),\, X_{H_2}(\pi)\big)$$

$$H_1 \sim H_2 \iff d(H_1, H_2) = 0, \qquad S_O := H/\!\sim$$

Proposition 2. *$\sim$ is the unique maximal monoidal congruence on H respecting all admissible observables.*

Proof. Existence. "Identical observables under all admissible protocols" generates $\sim$. The quotient exists.

Maximality. Any coarser relation retaining $H_1 \neq H_2$ despite $H_1 \sim H_2$ maintains a distinction with no invariant physical content. By A0, that distinction does not exist.

Sequential congruence. $H_1 \sim H_2$ implies $H_1 \circ H_3 \sim H_2 \circ H_3$: any protocol on $H_1 \circ H_3$ induces one on the first factor. If composition broke equivalence, H_3 would distinguish H_1 from H_2, contradicting $H_1 \sim H_2$.

Parallel congruence. For H_3 causally independent of H_1, H_2 (Step 2 primitive), test protocols on the parallel composition factor by causal independence. $H_1 \sim H_2$ implies agreement on the parallel composition.

Uniqueness. Two maximal monoidal congruences each imply the other by maximality. □

No circularity: causal independence is the Step 2 primitive; $\sim$ is defined afterward.
Remark 1. Real observers approximate this quotient limit with increasing fidelity as resources increase. This is the same idealization thermodynamics makes with reversible processes.

Step 6: Composition must respect the quotient

$$H_1 \sim H_2 \implies F(H_1) = F(H_2), \qquad F(H_1 \circ H_2) = F(H_1) \circ F(H_2)$$

Forced by A0 and Step 5.

Step 7: Locality — influence propagates at finite speed

Instantaneous influence would allow a protocol distinguishing histories whose only difference lies outside the observer's causal reach, violating Step 4. Therefore some finite $c > 0$ exists.

Step 8: Independent systems compose symmetrically; the monoid is cancellative

Two events are *causally disconnected* if no signal at speed $\leq c$ connects them.

Proposition 3. *For causally disconnected sectors A and B: $H_A \circ H_B \sim H_B \circ H_A$.*

Proof. A protocol distinguishing the orderings must carry information between sectors, requiring a signal at speed $\leq c$. No such signal exists. By Step 5 the orderings are identical. □

The quotient $(S_O, \oplus)$ restricted to causally disconnected sectors is therefore a commutative topological monoid—Hausdorff by the quotient metric, locally compact by Step 3, with jointly continuous composition and identity $[\emptyset]$ (the empty-history class).

Lemma 4 (Inversion — the monoid is cancellative)**.** *If $[H_1] \oplus [H_3] = [H_2] \oplus [H_3]$ in S_O for any causally independent H_3, then $[H_1] = [H_2]$.*

Proof. Suppose $H_1 \oplus H_3 \sim H_2 \oplus H_3$. For any admissible protocol π on H_1 alone, define the protocol π' on $H_1 \oplus H_3$ that executes π on the H_1 sector and does nothing on H_3 (possible by causal independence: protocols on causally independent sectors factor, so ignoring H_3 is an admissible restriction). Since $H_1 \oplus H_3 \sim H_2 \oplus H_3$, the outcome distributions of π' agree on both sides. The H_1-sector component of π' is exactly π, so π gives identical statistics on H_1 and H_2. Since π was arbitrary, $H_1 \sim H_2$. □

Corollary 5. *The Grothendieck group $\mathcal{G}(S_O)$ of the cancellative commutative monoid $(S_O, \oplus)$ is a locally compact Hausdorff abelian group, and S_O embeds into $\mathcal{G}(S_O)$ as the positive cone. All subsequent structure-theorem arguments apply to $\mathcal{G}(S_O)$.*

Proof. Elements of $\mathcal{G}(S_O)$ are formal differences $[a] - [b]$ with $[a] - [b] = [a'] - [b']$ iff $[a] \oplus [b'] = [a'] \oplus [b]$ (well-defined by cancellativity). The translation-invariant metric d

from Lemma 6 extends to $\mathcal{G}(S_O)$ by:

$$D\big([a]-[b],\ [a']-[b']\big) = d\big([a]\oplus[b'],\ [a']\oplus[b]\big)$$

This is well-defined (independent of representative by cancellativity), symmetric, satisfies the triangle inequality, and is translation-invariant by construction. Closed D-balls of radius r in $\mathcal{G}(S_O)$ are totally bounded: any $[a]-[b]$ with $D([a]-[b],0) \le r$ satisfies $d([a],[b]) \le r$, and bounded sets in (S_O, d) are totally bounded by Step 3 (Bekenstein bound). Completeness holds because S_O is complete (reachable states form a complete metric space under d). A complete totally bounded metric group is locally compact. Therefore $\mathcal{G}(S_O)$ is locally compact Hausdorff. □

Step 9: The state space is a finite-dimensional real vector space

$\mathcal{G}(S_O)$ is a connected (S_O is path-connected), locally compact, Hausdorff abelian group. It is Archimedean: if $x>0$ and $nx \le y$ for all $n \in \mathbb{N}$, then $x \sim 0$ by Step 5, contradicting $x>0$; by A0, permanently undetectable structure does not exist. By the structure theorem for connected locally compact abelian groups [9], any such group is isomorphic to $\mathbb{R}^n \times T^m$. The Archimedean property kills T^m: the torus contains elements of all finite orders ($n\theta = 0$ for some n, violating Archimedean for $\theta > 0$). Therefore:

$$S_O \hookrightarrow \mathcal{G}(S_O) \cong \mathbb{R}^n, \quad 2 \le n \le k$$

Step 10: The distinguishability measure is quadratic

The vector space $\mathbb{R}^n$ has already been derived in Step 9 from the group structure of S_O via the Grothendieck construction and the Pontryagin structure theorem. Only now, with $\mathbb{R}^n$ in hand, does the rotation argument apply. The logical order is: group structure $\Rightarrow \mathbb{R}^n \Rightarrow \mathrm{O}(n)$-invariance $\Rightarrow$ Euclidean norm. Linearity precedes the symmetry argument, not the reverse.

Any orthonormal basis of $\mathbb{R}^n$ is an equally valid description of the same physical state: basis choice is a convention with no physical content, by A0. The operational norm — the function that encodes physical distinguishability — must therefore be invariant under all orthonormal basis changes. These are exactly the orthogonal transformations $\mathrm{O}(n)$.

Lemma 6 (Induced norm is translation-invariant)**.** *The operational metric d is translation-invariant under $\oplus$: for all causally independent s, s', t in S_O,*

$$d(s \oplus t,\ s' \oplus t) = d(s, s').$$

Proof.

$$d(s \oplus t,\ s' \oplus t) = \sup_{\pi} D_{\mathrm{TV}}\big(X_{s\oplus t}(\pi),\ X_{s'\oplus t}(\pi)\big).$$

By causal independence (Step 2) and the monoidal congruence (Step 5), any admissible protocol π on $s \oplus t$ factors into an s-sector component π_s and a t-sector component π_t with no cross-talk (a joint protocol exploiting correlations between causally disconnected sectors would require a connecting signal, excluded by Step 7). The TV distance therefore decomposes: only the s-sector contributes to distinguishing $s \oplus t$ from $s' \oplus t$. Taking the supremum over all admissible protocols:

$$d(s \oplus t,\ s' \oplus t) = \sup_{\pi_s} D_{\mathrm{TV}}\big(X_s(\pi_s),\ X_{s'}(\pi_s)\big) = d(s, s'). \qquad \square$$

Corollary 7. *The operational metric d defines a norm on $\mathbb{R}^n$ by $\|x\| = d(x, 0)$. By A0, any orthonormal basis change is pure redescription (basis choice carries no physical content), so the norm must be invariant under all orthogonal transformations of $\mathbb{R}^n$. By Lemma 1, the unique such norm is Euclidean:*

$$\|x \oplus y\|^2 = \|x\|^2 + \|y\|^2.$$

Step 11: The state space has an inner product

The quadratic norm of Step 10 satisfies the parallelogram law:

$$\|x + y\|^2 + \|x - y\|^2 = 2\big(\|x\|^2 + \|y\|^2\big)$$

By the Jordan–von Neumann theorem [8], this is equivalent to the existence of an inner product:

$$\langle x, y\rangle = \tfrac{1}{4}\big(\|x + y\|^2 - \|x - y\|^2\big)$$

$S_O \cong \mathbb{R}^n$ with this inner product is a real finite-dimensional inner product space.

Step 12: The phase group exists and is U(1)

Empirical input. We take as an empirical fact about our universe that non-trivial free dynamics exist: physical systems can evolve continuously between interactions. This is consistent with all observations and is the content of "non-trivial time evolution." It plays the same role here that finite propagation speed plays at Step 7 — an empirical fact about the universe we inhabit that the framework incorporates.

Physical argument. Consider an isolated system undergoing no interactions. By Step 4, no interactions means no new physical knowledge for the observer. Therefore the observer's probability assignments for that system are constant under free evolution. Yet the physical state can change continuously — a precessing magnetic moment, a clock ticking, a spinning top. States and probability assignments are necessarily distinct objects.

Formal derivation. We use "probability assignment" informally here for the observer's predictive map F; its explicit form is derived in Step 16. The map:

$$\pi : S_O \to \text{Prob(outcomes)}$$

cannot be injective. If it were, every state change under free dynamics would change probability assignments; but free dynamics involve no interactions, so by Steps 4 and 5, probability assignments cannot change. Non-injectivity follows from the interaction-only epistemology applied to the case of no interaction.

Free dynamics must preserve the Euclidean norm of Step 10: if distinguishability could be altered by free evolution without interaction, information would be created or destroyed without physical cause, violating Steps 4 and 5. Continuous norm-preserving one-parameter groups on $\mathbb{R}^n$ are one-parameter subgroups of $\mathrm{O}(n)$, i.e. of the form $\exp(tA)$ where A is a skew-symmetric matrix. By the real canonical form for skew-symmetric matrices [2], A decomposes $\mathbb{R}^n$ into two-dimensional invariant planes (corresponding to purely imaginary eigenvalue pairs $\pm i\omega$) and fixed-point subspaces (zero eigenvalues). Within each invariant plane, the dynamics $\exp(tA)$ are continuous rotations at angular frequency ω. Since probability assignments are constant on dynamical orbits of isolated systems, and orbits within each plane are circles of constant radius, the probability map π depends only on the radius r within each dynamical plane. The fiber of π within each dynamical plane is $S^1 \cong \mathrm{U}(1)$.

Characterization. The phase group P — the group of norm-preserving transformations preserving all probability assignments — is compact and connected (the fiber S^1 is already compact and connected; a transformation leaving it becomes detectable), abelian (global phases commute with all physical operations), and minimal (Step 3 forbids redundant structure). A compact connected abelian Lie group is $T^k = \mathrm{U}(1)^k$. Minimality forces $k = 1$:

$$P = \mathrm{U}(1), \qquad \psi = r \cdot e^{i\theta}$$

Step 13: The scalar field is $\mathbb{C}$

The scalar field $\mathbb{K}$ must contain U(1) as a multiplicative subgroup (Step 12), support additive composition, be commutative (Step 8), associative (Step 6), and minimal (A0 and Step 3). By Hurwitz's theorem [7], the finite-dimensional normed division algebras over $\mathbb{R}$ are exactly $\mathbb{R}$, $\mathbb{C}$, $\mathbb{H}$, $\mathbb{O}$:

- $\mathbb{R}$: unit group $\mathbb{Z}_2$, does not contain U(1). Excluded.
- $\mathbb{H}$: non-commutative. Excluded by Step 8.
- $\mathbb{O}$: non-associative. Excluded by Step 6.
- $\mathbb{C}$: satisfies all requirements. Unique solution.

Therefore $\mathbb{K} = \mathbb{C}$.

Step 14: The state space is a complex Hilbert space

Complexification. The U(1) action from Step 12 defines a linear map $J : \mathbb{R}^n \to \mathbb{R}^n$ corresponding to the 90° rotation on each dynamical two-dimensional plane, with $J^2 = -I$. This is a complex structure on $\mathbb{R}^n$ in the standard sense (a linear automorphism squaring to $-I$), requiring n to be even; write $n = 2m$. The complexification $\mathbb{R}^{2m} \cong \mathbb{C}^m$ is canonical. The real inner product from Step 11 and J together define a Hermitian inner product:

$$\langle x, y\rangle_{\mathbb{C}} = \langle x, y\rangle_{\mathbb{R}} + i\langle Jx, y\rangle_{\mathbb{R}}$$

which is sesquilinear, conjugate-symmetric ($\langle y, x\rangle_{\mathbb{C}} = \overline{\langle x, y\rangle_{\mathbb{C}}}$), and positive definite. $H_O \cong \mathbb{C}^m$ with this Hermitian inner product is a finite-dimensional complex Hilbert space with $\dim_{\mathbb{C}}(H_O) = m \leq k/2$.

Local tomography. By Step 5, composite states are equivalent iff all admissible protocols agree. Admissible protocols on composite systems factor through local operations plus correlations by causal independence (Step 2). Therefore $\{A \otimes B : A \in \mathcal{L}(H_A),\, B \in \mathcal{L}(H_B)\}$ must generate all of $\mathcal{L}(H_{AB})$: if it did not, two composite states could agree on all local-plus-correlation statistics while differing on some joint observable, contradicting Step 5. Local tomography is derived from the monoidal congruence.

Steps 14–15 constraint system. Local tomography (derived here from Step 5 independently of any tensor assumption) and tensor-product composition (Step 15) form a coupled constraint system. The logical direction is strictly one-way: Step 5 $\Rightarrow$ local tomography $\Rightarrow$ tensor product $\Rightarrow$ J_3^8 exclusion. The coupling is a fixed point of compositional consistency, not a circularity.

Exclusion of J_3^8. Local tomography plus Step 15 tensor structure places the framework in the GPT setting. The Barnum–Wilce theorem [5] establishes that no locally tomographic homogeneous self-dual theory containing a qubit admits J_3^8. Both hypotheses are earned: local tomography from Step 5; the qubit from dim ≥ 2 in Step 3.

Step 15: Composite systems compose by tensor product

Theorem 8. *For causally independent systems A and B with Hilbert spaces H_A and H_B: $H_{AB} \cong H_A \otimes H_B$.*

Proof. Define $\iota : S_A \times S_B \to H_{AB}$ by $\iota(a, b) = [a \otimes b]$.

Bilinear. S_A and S_B are real vector spaces (Step 9). Protocols linear in sector A with fixed B map to linear combinations in H_{AB} by Step 5. Linearity in each argument follows independently, so ι is bilinear.

Spanning. By local tomography (Step 14), every composite state is fully determined by its overlaps with product states. The product states $\{\iota(a, b)\}$ therefore span H_{AB}.

Universal property. For any Hilbert space K and bounded bilinear $f : S_A \times S_B \to K$ consistent with independent protocols, define $\tilde{f}(\iota(a, b)) = f(a, b)$ and extend by linearity. This is well-defined (product states span H_{AB}), unique (two linear extensions agreeing on a spanning set are equal), and satisfies $f = \tilde{f} \circ \iota$.

Minimality. A space larger than $H_A \otimes H_B$ contains vectors unreachable by any bilinear protocol; by A0 and Step 5 those vectors have no physical existence. A space smaller fails to represent some admissible bilinear protocol, contradicting Steps 4 and 5. Therefore $H_{AB} \cong H_A \otimes H_B$, unique up to isomorphism. □

Entanglement. States not factorable into product states emerge immediately:

$$\exists\, \psi \in H_{AB} : \psi \notin \overline{\mathrm{conv}}\big\{\iota(a,b) : a \in H_A,\, b \in H_B\big\}$$

Entanglement is joint invariant structure that cannot be factored into individual parts.

Step 16: The Born rule

We use the term "probability assignment" formally from here: it is the map $\pi : S_O \to \mathrm{Prob(outcomes)}$ introduced in Step 12, which we now derive explicitly. Probabilities must satisfy:

- *Phase invariance*: $P(e^{i\theta}\psi) = P(\psi)$ for all θ — phases live in the U(1) fiber, which by construction preserves all probability assignments.
- *Multiplicativity*: $P(\psi \otimes \phi) = P(\psi)P(\phi)$ for independent systems — this is what independence means.
- *Continuity* under refinement.

The unique solution is:

$$P(\psi) = \langle \psi, \psi \rangle = |\psi|^2$$

Non-contextuality and what Gleason actually does. The Step 5 equivalence relation defines physical states as equivalence classes under *all* admissible protocols. Non-contextuality is therefore a consequence of what "state" means, not a separately derived theorem. Contextuality would require the same physical state to yield different distributions in different measurement contexts. But if two contexts yield different distributions, those contexts are admissible protocols that are distinguishing something — by Step 5, the situations are therefore *different states*, not the same state behaving contextually. Contextuality is impossible given the definition of state identity. This is not circular: it is the direct meaning of the quotient construction applied to the case of measurement context.

What Gleason's theorem then contributes is non-trivial and independent: it characterizes the unique *frame function* on the orthomodular lattice $\mathcal{L}(H)$ of closed subspaces of H. A frame function is a map from subspaces to $[0,1]$ summing to 1 on any orthonormal basis

— exactly the structure a probability assignment must have over complete measurement contexts. The orthomodular lattice follows from the quadratic norm (Step 10) and inner product (Step 11). Gleason proves that the unique such function is $\mathrm{Tr}(\rho\, P_E)$. This is a theorem about the lattice structure derived from prior steps, not an imported assumption:

$$P(E) = \mathrm{Tr}(\rho\, P_E)$$

for unique density operator ρ with $\mathrm{Tr}\,\rho = 1$.

The qubit closed internally. $\dim(H_O) \geq 2$ (Step 3). Any system S observed by O has joint space $H_O \otimes H_S$ with $\dim \geq 2 \cdot \dim(H_S)$. For any nontrivial S, $\dim(\text{joint}) \geq 4 \geq 3$. Gleason applies to the joint space. The Born rule for H_S follows by partial trace:

$$P(E_S) = \mathrm{Tr}_O\big(\rho_{OS} \cdot I_O \otimes P_{E_S}\big)$$

The observer is always present (Step 1). The isolated $\dim = 2$ case never arises within the framework.

Step 17: Dynamics are unitary

Free evolution must preserve the invariant structure — the Euclidean norm of Step 10. If distinguishability could be destroyed by free dynamics without interaction, information would be lost in violation of Steps 4 and 5. The norm-preserving maps are isometries.

By Wigner's theorem [12], every isometry of a complex Hilbert space is unitary or antiunitary. Continuity in d excludes antiunitary one-parameter families for time evolution. By Stone's theorem [11]:

$$U(t) = e^{-iHt}, \quad H \text{ self-adjoint}$$

This formalizes what Step 12 established physically: the continuous norm-preserving dynamics forced by Steps 4 and 5 are exactly unitary evolution.

Final Statement

In one sentence. If existence is constituted by interaction, then the structure of physics is exactly and uniquely what a finite observer inside the world must find — and that structure is quantum mechanics.

Theorem 9 (Main)**.** *In any universe containing finite, local, relational, memory-bearing, physically situated observers whose existence is constituted by interaction (Steps 1–5) and who therefore recognize A0 as the minimal commitment consistent with that existence, the physics accessible to those observers is, in the finite-dimensional setting, uniquely:*

- ***State space****: finite-dimensional complex Hilbert space H*
- ***Composition****: tensor product $H_A \otimes H_B$*
- ***Probabilities****: Born rule $P(E) = \mathrm{Tr}(\rho\, P_E)$*
- ***Dynamics****: unitary $U(t) = e^{-iHt}$*

Alternatives excluded at each step.

- Classical probability: no phase group. Fails at Step 12.
- Real Hilbert space: no continuous phase structure. Fails at Steps 11–12.
- Quaternionic QM: non-commutative. Fails at Steps 8 and 13.
- Super-quantum theories: non-Euclidean norm. Fail at Step 10.
- Exceptional Jordan algebra J_3^8: fails local tomography. Fails at Step 14.

Honest Boundary

The derivation is finite-dimensional throughout. Infinite-dimensional Hilbert space, QFT, and GR are outside scope and explicitly not claimed.

A0 is recognition, not theorem — the one commitment every reader has already made by being a finite physical system reading this. The Steps 14–15 coupling is a fixed point of compositional consistency with explicit logical direction: Step 5 $\Rightarrow$ local tomography $\Rightarrow$ tensor product $\Rightarrow$ J_3^8 exclusion.

Special relativity is consistent with and illuminates the geometry of Step 8. The abelian composition result follows from finite-c and observer constraints alone.
Open Gap 1. The infinite-dimensional limit: extending the derivation to separable infinite-dimensional Hilbert spaces, and ultimately to the QFT setting, is the natural next problem. The finite-dimensional theorem as stated is complete.

Acknowledgments

The author thanks the adversarial review process, without which several gaps would have remained hidden.

References

[1] G. Chiribella, G.M. D'Ariano, and P. Perinotti, "Informational derivation of quantum theory," *Phys. Rev. A* **84**, 012311 (2011).

[2] B.C. Hall, *Lie Groups, Lie Algebras, and Representations: An Elementary Introduction*, 2nd ed. (Springer, New York, 2015). [For the real canonical form of skew-symmetric matrices and decomposition of one-parameter subgroups of O(n) into two-dimensional invariant planes, see Ch. 2–3.]

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

[4] J.D. Bekenstein, "Black holes and entropy," *Phys. Rev. D* **7**, 2333 (1973).

[5] H. Barnum and A. Wilce, "Local tomography and the Jordan structure of quantum theory," *Found. Phys.* **44**, 192–212 (2014).

[6] A.M. Gleason, "Measures on the closed subspaces of a Hilbert space," *J. Math. Mech.* **6**, 885–893 (1957).

[7] A. Hurwitz, "Über die Komposition der quadratischen Formen," *Math. Ann.* **88**, 1–25 (1923).

[8] P. Jordan and J. von Neumann, "On inner products in linear, metric spaces," *Ann. Math.* **36**, 719–723 (1935).

[9] L.S. Pontryagin, *Topological Groups*, 2nd ed. (Gordon and Breach, New York, 1966).

[10] G. Sobczyk, "Killing spinors in quantum mechanics," *Adv. Appl. Clifford Algebras* **17**, 811–820 (2007). [For O(n)-invariant norms see also F. Bohnenblust, "Convex regions and projections in Minkowski spaces," *Ann. Math.* **39**, 301–308 (1938).]

[11] M.H. Stone, "On one-parameter unitary groups in Hilbert space," *Ann. Math.* **33**, 643–648 (1932).

[12] E.P. Wigner, "Gruppentheorie und ihre Anwendung auf die Quantenmechanik der Atomspektren" (Vieweg, Braunschweig, 1931). [For Wigner's theorem see also V. Bargmann, "Note on Wigner's theorem on symmetry operations," *J. Math. Phys.* **5**, 862 (1964).]