Why This Physics? Quantum Mechanics as the Architecture of Fact-Production

Keith Taylor

VERSF Theoretical Physics Program

Quantum mechanics isn't one possible physics among many—it's the unique structural framework compatible with a universe where facts can be recorded. This paper proves it.

Abstract

We show that quantum mechanics is not one admissible theory among many, but the unique structural framework compatible with physical admissibility in a fact-producing universe. Starting from two minimal constraints—finite operational distinguishability and irreversible commitment—we derive the structural core of quantum mechanics without presupposing Hilbert spaces, probability measures, or dynamical laws. We establish that infinite distinguishability is incompatible with irreversible facts, forcing any fact-producing physics to operate within finite resolution bounds (Tier I: forced by admissibility). From these constraints plus operational closure principles—tomographic locality and maximal reversibility—we demonstrate that admissible reversible evolution must be unitary and Hamiltonian-generated, that measurement emerges as the minimal irreversible extension of unitary dynamics, and that the Born rule is the unique admissible probability assignment (Tier II: selection among admissible theories). Entanglement and Bell correlations follow from global commitment structure, while first-order locality forces Clifford algebra and Dirac dynamics. These results are structural and interpretation-independent, relocating quantum foundations from postulates to necessity.

Plain Language Overview

Note: This article is written in two layers. The main text is technical and theorem-driven; the "In Plain Language" boxes give an intuitive explanation of what each result means. Feel free to read at whichever level suits you—or both.

What this paper shows, in everyday terms:

Quantum mechanics—the physics of atoms, electrons, and light—is famously strange. Particles can be in two places at once, measurements seem to disturb what they measure, and distant particles can be mysteriously correlated. For a century, physicists have debated what this strangeness *means* and whether we simply have to accept these rules as brute facts about nature.

This paper takes a different approach. Instead of asking "why is quantum mechanics so weird?", we ask: "what kind of physics could possibly allow *facts* to exist?"

Think about what a "fact" requires. When you measure something—say, whether a light is on or off—you get a definite answer that stays definite. The measurement creates a *record* that can't be undone. This seems obvious, but it turns out to place severe constraints on what kind of physics is possible.

We establish that if facts can exist at all, then:

- There must be a limit to how precisely you can distinguish things (you can't have infinite resolution)
- Between measurements, physics must evolve in a very specific way (unitarily, governed by energy)
- Probabilities must follow a specific rule (the Born rule, with its squared amplitudes)
- Particles can be "entangled" in ways that create correlations without sending signals
- Relativistic particles must have "spin" and obey the Dirac equation

. **If facts can be recorded at all, quantum structure is inevitable.**

In other words, quantum mechanics isn't one possible physics among many—it's the unique *structural framework* compatible with a universe where facts exist and records persist. The strangeness isn't optional; it's mathematically inevitable.

The title calls quantum mechanics an "admissibility fixed point." A *fixed point* is the structure you inevitably arrive at after applying constraints again and again—nothing else survives. We demonstrate that once you demand facts be possible, you can keep asking "what does that require?" and every answer points to quantum mechanics. It's where the logic terminates.

Contents:

1. Introduction: The Problem of Quantum Foundations

Quantum mechanics is among the most successful theories in the history of science. Its empirical accuracy is unmatched, and its mathematical formalism has been confirmed across an extraordinary range of physical regimes. Yet despite this success, the foundational status of the theory remains unsettled. The core mathematical structures of quantum mechanics—Hilbert spaces, unitary evolution, the Born rule, and the projection postulate—are introduced largely by postulate rather than necessity. The theory specifies how physical systems evolve and how measurement outcomes are calculated, but not *why* this particular structure is required for a physically meaningful description of the world.

This lack of structural grounding has given rise to a long-standing foundational impasse. Interpretational programs debate the ontological meaning of the quantum state while holding the formalism fixed. Reconstruction approaches attempt to re-derive quantum theory from informational or operational axioms [1–5], yet typically presuppose outcome finality, probabilistic structure, or reversible dynamics without explaining why these features must coexist. Decoherence-based accounts successfully explain the suppression of interference in open systems [6] but do not, by themselves, account for the emergence of definite outcomes or irreversible records. As a result, the measurement problem persists—not as a technical inconsistency within the formalism, but as a symptom of a deeper structural question left unanswered.

The central difficulty may be stated succinctly: quantum mechanics presupposes the existence of irreversible facts while simultaneously employing a fundamentally reversible mathematical core. Measurements yield definite outcomes; records persist; entropy increases; time exhibits a directed character. Yet the dominant dynamical structure of the theory—unitary evolution on a Hilbert space—is perfectly reversible.

In this work, we propose a shift in perspective. Rather than asking how quantum mechanics should be interpreted, we ask a logically prior question: *what constraints must any physically admissible theory satisfy in order to produce irreversible facts, persistent records, and a meaningful arrow of time?*

1.1 Relation to Existing Reconstruction Programs

Several programs have sought to derive quantum mechanics from operational or information-theoretic principles. Hardy's approach [1] derives quantum theory from five "reasonable axioms" including probabilistic structure and a simplicity postulate. Chiribella, D'Ariano, and Perinotti [2] obtain quantum theory from operational primitives assuming the "purification postulate." Masanes and Müller [3] derive the formalism from information-processing constraints. Dakić and Brukner [4] emphasize the role of continuous reversibility.

The present work differs from these programs in a crucial respect: we do not assume probabilistic structure, reversible dynamics, or Hilbert space at the outset. Instead, we begin with a more primitive question—what is required for *facts* to exist?—and show that probabilistic structure, unitarity, and Hilbert space representation emerge as necessary consequences of admissibility constraints. The framework thus operates at a logically prior level, explaining *why* the assumptions of other reconstruction programs are themselves necessary.

Unlike prior reconstructions, which assume measurement and probability to derive kinematics, we derive the kinematic and probabilistic structure of quantum mechanics from the existence of irreversible facts themselves.

The key distinction: other programs ask "what must a theory with measurements and probabilities look like?" while we ask the logically prior question "what must a theory have to produce measurements and probabilities at all?"

Specifically, axioms that other programs assume become theorems here:

Other programs assume	This work derives
Probabilistic structure	Theorem 7.1 (Born rule from admissibility)
Hilbert space	Lemma 5.1 (from Jordan algebra classification)
Unitary dynamics	Theorem 5.1 (from distinguishability preservation)
Continuous reversibility	Stone's theorem applied to admissible evolution
Purification (CDP)	Follows from Hilbert space structure

What we do assume explicitly—the existence of facts (A1), finite resources (A2), and measurement realizability (A3)—are preconditions for any empirical physics, not specific to quantum theory. See Appendix C for detailed comparison.

1.2 Summary of Results

The results of this paper fall into two tiers with different logical status:

Tier I: Forced by Admissibility

Any fact-producing physical theory must exhibit:

- Finite operational distinguishability (Theorem 3.2)
- Irreversible commitment as the mechanism of fact-production
- Reversible dynamics that preserve distinguishability
- Measurement as many-to-one, entropy-increasing process

These constraints are *necessary consequences* of admissibility. They exclude infinite-precision classical mechanics and force a quantum-like reversible sector. No additional assumptions are required.

Tier II: Selection Among Admissible Theories

Additional operational principles select complex Hilbert-space quantum mechanics from among the admissible class. Crucially, these principles are not independent postulates but are *derivable* from minimal assumptions about experimental capabilities (see Appendix E):

- Convexity (R1) follows from classical control capability (Lemma E.1)
- Tomographic locality (R4) follows from universal composite controllability (Theorem E.5)
- Purification (R5) follows from reversible embeddability (Theorem E.3)

Together with R2–R3 (directly forced by admissibility), these yield uniquely: unitary dynamics (Theorem 5.1), Hamiltonian generation (Theorem 5.2), the Born rule (Theorem 7.1), entanglement structure (Section 8), and Dirac dynamics under relativistic extension (Section 9).

The Central Claim

Admissibility carves out a narrow class of possible physical theories. Within that class, minimal operational completeness—the requirement that a theory describe all experimentally accessible degrees of freedom—uniquely selects complex quantum mechanics. The assumptions required for Tier II (classical control, reversible embeddability, universal controllability) are not quantum-specific axioms but capabilities any laboratory doing physics must have.

---

2. Operational Framework and Physical Admissibility

In Plain Language: This section sets up the basic vocabulary we need. We define what it means for a physical process to be "admissible" (actually doable with finite resources), and we introduce two key concepts: "distinguishability" (can you tell two things apart?) and "commitment" (has a definite fact been established?). These aren't quantum mechanical ideas—they're more basic than any particular physics. They're about what it takes for *any* physics to produce facts.

Everyday example of admissibility: A digital thermometer displays temperature to one decimal place—say, 98.6°F. The *mathematical* temperature might be 98.6417293...°F, but the *operational* temperature is just 98.6°F. That's all the device can distinguish. Similarly, a digital photo has finite pixel resolution; a detector has a threshold below which it won't trigger; an electronic signal has finite bit depth. In every case, the physical system has a limit to how fine a distinction it can make. That limit is what we mean by "admissibility"—the boundary between what's mathematically describable and what's physically accessible.

In this section we introduce the minimal operational framework required to formalize the notion of physical admissibility. The goal is not to assume quantum mechanics or any particular dynamical law, but to specify the structural conditions under which physical distinctions, processes, and records can be meaningfully discussed.

Definition (Operational Fact). In this work, a *fact* means a physically realized, stable, recordable distinction accessible to finite physical procedures. This operational notion is the only sense in which facts enter empirical physics. Claims about ontological facts beyond operational access—distinctions that exist "in themselves" but cannot be detected, recorded, or referenced by any finite procedure—are metaphysical and lie outside the scope of physical theory.

Scope clarification. All results in this paper concern operational facts. When we say "facts require finite distinguishability," this is to be understood in the operational sense relevant to physics. We make no claims about metaphysical facts that transcend operational access; such claims are neither supported nor refuted by our analysis.

Throughout this work, we distinguish sharply between mathematical possibility and physical admissibility. Mathematical models may permit infinite precision, unbounded state spaces, or perfectly reversible dynamics without contradiction. Physical admissibility, by contrast, concerns which distinctions and transformations can be realized by finite physical procedures. The admissibility layer therefore functions as a constraint filter on formal theories, determining which structures are capable of supporting facts, records, and irreversible outcomes.

Neutrality on scientific realism. The operational stance is methodological, not metaphysical. The results are neutral on whether physics describes "reality itself" or only operational procedures. A scientific realist can read the framework as constraining what any correct description of reality must look like at the operational level. An instrumentalist can read it as constraining predictively adequate theories. The admissibility constraints apply equally under either interpretation: any physically meaningful theory—whether interpreted realistically or operationally—must satisfy them to be empirically testable.

2.1 Operational State Spaces

We model the set of physically preparable states of a bounded system by a convex state space $\mathcal{S}$. Elements of $\mathcal{S}$ represent equivalence classes of physical preparations that cannot be distinguished within the operational limits of the system. Convex combinations represent classical mixing of preparation procedures.

We assume that operational distinguishability between states is quantified by a function

$D : \mathcal{S} \times \mathcal{S} \to [0,1]$,

with the interpretation that $D(\rho,\sigma)$ measures the optimal probability of distinguishing ρ from σ using admissible measurement procedures. No particular functional form is assumed at this stage. The only required properties are:

(i) *Non-negativity and symmetry*: $D(\rho,\sigma) = D(\sigma,\rho) \geq 0$, with equality if and only if ρ and σ are operationally indistinguishable.

(ii) *Contractivity under admissible processes*: for any admissible process Φ,

$D(\Phi(\rho), \Phi(\sigma)) \leq D(\rho, \sigma)$.

This expresses the principle that physical processes cannot increase distinguishability beyond what is operationally accessible.

Examples of such distinguishability measures include the trace distance $D(\rho,\sigma) = \frac{1}{2}\|\rho - \sigma\|_1$ and the Helstrom optimal discrimination probability in quantum theory [7], but the framework does not assume Hilbert space structure *a priori*.

2.2 Admissible Processes

An admissible process is a map

$\Phi : \mathcal{S} \to \mathcal{S}$

that can be physically implemented using finite resources. By resources we mean finite energy, finite time, finite spatial extent, finite memory, and finite resolution. The precise accounting of these resources is left implicit; what matters is that admissibility excludes procedures that require infinite precision, infinite memory, or unbounded refinement.

Admissible processes include reversible transformations, irreversible transformations, and measurement procedures. They form a closed class under composition: if Φ_1 and Φ_2 are admissible, then $\Phi_2 \circ \Phi_1$ is admissible.

We emphasize that admissibility is an operational notion rather than a metaphysical one. A process may be mathematically well-defined yet physically inadmissible if its implementation requires unbounded control or resolution.

2.3 Finite Distinguishability

We now introduce the first core admissibility constraint.

Definition 2.1 (Finite Distinguishability). A bounded physical system exhibits *finite distinguishability* if there exists a finite bound N such that no admissible procedure can reliably distinguish more than N mutually distinguishable states within that system.

In Plain Language: Imagine you have a dial that can point anywhere on a circle. Mathematically, there are infinitely many positions. But physically, with any real measuring device, you can only reliably tell apart a finite number of positions—maybe 100, maybe a million, but not infinity. That's finite distinguishability. No matter how good your instruments get, there's always a limit to how fine a distinction you can make with finite resources.

Remark on scaling. The bound N depends on system size, available energy, and operational resources. For quantum systems, N scales with Hilbert space dimension, which in turn scales exponentially with particle number for composite systems. The information-theoretic content of a bounded region is constrained by holographic bounds ($S \leq A/4$ in Planck units) [22, 23], suggesting that N is ultimately bounded by geometry. What matters for the present analysis is not the precise value of N but its finiteness under any fixed resource budget.

Equivalently, the operational state space admits a maximum effective cardinality under admissible discrimination. This constraint does not require the mathematical state space to be finite or discrete. Continuous state spaces are permitted, provided that operational access to distinctions terminates at a finite resolution.

Finite distinguishability is a capacity constraint rather than a statement about ontology. It asserts that, under fixed finite resources, there exists a terminal refinement scale beyond which distinctions are not physically accessible.

2.4 Irreversible Commitment

The second core admissibility constraint concerns the physical realization of facts.

Definition 2.2 (Irreversible Commitment). An admissible process Φ exhibits *irreversible commitment* if there exist distinct states $\rho_1 \neq \rho_2$ in $\mathcal{S}$ such that

$\Phi(\rho_1) = \Phi(\rho_2)$,

and no admissible process Ψ exists satisfying $\Psi \circ \Phi(\rho_i) = \rho_i$ for all such inputs.

In Plain Language: When you take a photograph, light from a scene is captured on film or a sensor. The original 3D scene had vastly more information than the 2D image preserves. Once the photo is taken, you can't reconstruct everything about the original scene from the photo alone—information has been irreversibly lost. That's commitment: multiple different starting points lead to the same outcome, and you can't undo it to recover which starting point you had. Every measurement, every record, every memory works this way.

Irreversible commitment is defined relative to the admissible domain. It does not assert that information is destroyed in an absolute or metaphysical sense, but that distinctions become inaccessible to all admissible recovery procedures. Measurement outcomes, persistent records, and memory formation are paradigmatic examples of irreversible commitment.

This definition excludes mere practical irreversibility. A process is irreversibly committing only if recovery is inadmissible *in principle*, not merely difficult in practice.

2.5 Admissibility as a Constraint Layer

The admissibility framework introduced here is deliberately minimal. It does not assume quantum mechanics, classical mechanics, or any specific dynamical law. It introduces no stochastic postulates and no interpretational commitments. Its sole function is to specify which distinctions and transformations can be physically realized.

In the sections that follow, we demonstrate that these constraints are already sufficient to exclude broad classes of candidate physical frameworks and to force the mathematical structure of quantum mechanics.

2.6 A Concrete Example: The Digital Camera Pixel

To make the admissibility framework vivid, consider a single pixel in a digital camera sensor. This everyday example illustrates every element of the framework: finite distinguishability, irreversible commitment, entropy increase, and the boundary between reversible and irreversible physics.

The setup. A camera pixel consists of a photosensitive region that converts incoming photons into electrical charge, an analog charge storage well, an analog-to-digital converter (ADC), and a memory register that stores a discrete numerical value (say, 0–255).

Before the shutter closes, light from the scene strikes the pixel. Photons arrive at random times and deposit energy, gradually building up charge. At this stage, the system is still reversible in principle: the charge is an analog quantity, and tiny differences in photon arrival history still exist in the microscopic degrees of freedom. This is the *pre-commitment phase*.

Finite distinguishability. The pixel does not record an arbitrary real number of charge. The ADC divides the possible charge range into 256 bins; any charge value within a given bin produces the same digital output. Mathematically, infinitely many distinct microscopic charge configurations exist. Operationally, only 256 distinguishable outcomes are accessible.

This limitation is not merely technological. Even with better electronics, noise, thermal fluctuations, and finite energy place a bound on how finely charge can be resolved. No admissible physical procedure can extract infinite precision from a bounded pixel in finite time. This is exactly what finite operational distinguishability means.

Irreversible commitment. When the shutter closes, the camera performs analog-to-digital conversion and writes a number to memory. At this moment, many different microscopic charge configurations map to the same digital value. The detailed photon arrival history is discarded. No admissible process can reconstruct which microscopic configuration occurred.

Two physically distinct pre-measurement states—slightly more photons early versus slightly more photons late, or different microscopic noise realizations—now produce the same recorded pixel value. Once the number is written, the fact is created. The pixel is now "value = 137," full stop.

Entropy increase. Before commitment, the pixel's microscopic state contained fine-grained information about photon arrivals, thermal noise, and charge distribution. After commitment, all that information is compressed into one of 256 labels. The number of physically distinguishable states has decreased; information about the past has been irreversibly erased. This entropy increase follows directly from finite distinguishability and irreversible many-to-one mapping. Entropy increases because facts require erasure.

Loss of reversibility. Could we undo the measurement? You can delete the photo file, power off the camera, or reset the memory. But none of these operations recover the exact charge configuration, the photon arrival times, or the original light field. The system has crossed a one-way boundary. Reversibility fails not because physics is approximate, but because the distinction has been committed into a finite record and admissible operations cannot access the erased degrees of freedom.

The lesson. This single pixel illustrates the entire admissibility framework—and nothing in this example is "quantum weirdness." It is ordinary, everyday physics. Quantum mechanics generalizes this structure to all physical systems, including microscopic ones. What we call "wavefunction collapse" is not exotic; it is the same kind of irreversible commitment that happens every time a camera records a pixel value.

If you demanded infinite pixel resolution, perfect reversibility, and zero information loss, then no photograph could ever exist as a fact. The image would remain a perpetually reversible physical process, never producing a stable record. The admissibility constraints are not quantum-specific—they are *fact-specific*. Quantum mechanics is simply the theory that respects these constraints universally.

3. Finite Distinguishability and Irreversible Commitment

In Plain Language: This section contains the paper's first major result: a proof that you can't have both infinite precision *and* real facts. If you could always zoom in further and see finer details, then nothing would ever be truly settled—any apparent "fact" could be undone by accessing those finer details. So if facts exist (and they do—you're reading one right now), then there must be a bottom level, a finest grain, beyond which distinctions can't be made. This isn't a feature of quantum mechanics specifically; it's a requirement for *any* physics that produces facts.

Important clarification: "Infinite distinguishability" doesn't mean "the math uses real numbers" (it does, and that's fine). It means the universe would let you *physically access* arbitrarily fine distinctions as operational facts. The math can be continuous; what matters is whether nature provides a physical procedure to resolve every mathematical distinction. We argue it cannot—not if facts are to exist.

In this section we establish the central structural result of the admissibility framework: irreversible commitment is impossible in any physical theory that permits infinite operational distinguishability. This result functions as a no-go theorem. It does not depend on quantum mechanics, probabilistic assumptions, or dynamical laws.

3.1 The Problem of Commitment

Physical theories are required to account for the existence of facts. Measurements yield definite outcomes, records persist, memories form, and entropy increases. These phenomena share a common structural feature: multiple prior possibilities resolve into a single outcome that cannot be undone within the operational domain. We refer to this process as irreversible commitment.

Many mathematical models of physics permit apparent irreversibility only as a result of coarse-graining or incomplete description. If finer degrees of freedom remain physically accessible, then any apparent loss of information can, in principle, be reversed. The key question is therefore not whether a process appears irreversible at a given descriptive level, but whether the distinctions erased by that process are inaccessible to all admissible refinements.

3.2 Refinement Lemma

We first formalize the notion of refinement that underlies distinguishability.

Lemma 3.1 (Refinement Lemma). Let ρ_1 and ρ_2 be two operationally distinguishable states in the admissible state space $\mathcal{S}$. Then there exists an admissible measurement or refinement variable Φ such that $\Phi(\rho_1) \neq \Phi(\rho_2)$.

Proof. By definition, operational distinguishability means that there exists an admissible procedure that reliably yields different outcomes when applied to ρ_1 and ρ_2. Let Φ denote the outcome variable of this procedure. Since the procedure is admissible, Φ corresponds to a physically accessible refinement that separates the two states. ■

This lemma makes explicit that distinguishability commits one to the existence of physically accessible degrees of freedom that encode the distinction.

3.3 No-Go Theorem for Infinite Distinguishability

We now state the main result. First, we make precise what infinite distinguishability means operationally.

Definition 3.1 (Infinite Distinguishability). A system exhibits *infinite distinguishability* if for every $\varepsilon > 0$ and every pair of distinct states $\rho_1 \neq \rho_2$, there exists an admissible procedure that distinguishes them with error probability less than ε. Equivalently, the operational distinguishability function $D(\rho_1, \rho_2)$ can be made arbitrarily close to 1 for any distinct pair using admissible procedures.

This definition does not presuppose recoverability—it only asserts arbitrarily good discrimination. The theorem shows that such discrimination power implies recoverability, which then excludes commitment.

Remark on ontology vs. operation. A potential objection: "What if infinitely many states exist ontologically but only finitely many are operationally accessible?" This objection actually supports the framework. We make no claims about what exists "in itself"—only about what can be operationally accessed. If operational access is bounded while ontology is richer, then operational physics (the subject of empirical science) satisfies finite distinguishability. The theorem constrains operational physics, which is what experiments test. Claims about inaccessible ontology are metaphysics, not physics. See Appendix B.1 for extended discussion.

Theorem 3.2 (No-Go for Infinite Distinguishability). If an admissible state space permits infinite distinguishability, then no admissible process can exhibit irreversible commitment.

The Intuition: If you can always look closer and see more detail, then any information you thought was "lost" is actually still there at a finer level. Nothing is ever truly erased; it's just hidden. But facts require genuine erasure—genuine loss of alternatives. So infinite resolution and genuine facts are incompatible.

Proof. Assume, for contradiction, that the state space permits infinite distinguishability (Definition 3.1) and that there exists an admissible process Φ exhibiting irreversible commitment. Then there exist distinct states $\rho_1 \neq \rho_2$ such that $\Phi(\rho_1) = \Phi(\rho_2) = \rho^*$.

Step 1: Discrimination implies information encoding. By infinite distinguishability, for any $\varepsilon > 0$ there exists an admissible procedure M_ε that distinguishes ρ_1 from ρ_2 with error $< \varepsilon$. This procedure extracts information about which state was prepared. Physically, this information must be encoded in degrees of freedom accessible to admissible operations.

Step 2: Accessible information survives admissible transformations. Let the output of M_ε be a classical record r_ε indicating whether the input was ρ_1 or ρ_2 (with error $< \varepsilon$). Consider the composite procedure: first apply M_ε to create r_ε, then apply Φ. The record r_ε is created by an admissible procedure and is therefore accessible to admissible operations.

Step 3: The paradox. After Φ acts, we have $\Phi(\rho_1) = \Phi(\rho_2) = \rho^*$. But the record r_ε, created before Φ was applied, still exists and still encodes (with error $< \varepsilon$) whether the original state was ρ_1 or ρ_2. Therefore, even though $\Phi(\rho_1) = \Phi(\rho_2)$, the distinction between having started from ρ_1 versus ρ_2 remains accessible via r_ε.

Step 4: Contradiction. The composite system (output state ρ^* + record r_ε) retains the ability to distinguish original preparations. An admissible recovery procedure exists: read r_ε. This contradicts the assumption that Φ exhibits irreversible commitment, which requires that no admissible procedure can recover the ρ_1–ρ_2 distinction after Φ.

Since ε was arbitrary, the argument holds at every precision level. Therefore, under infinite distinguishability, no admissible process can irreversibly commit distinctions. ■

Remark on avoiding circularity. The key logical structure is: (1) infinite distinguishability is defined purely in terms of discrimination ability, not recoverability; (2) the proof shows that discrimination ability implies the existence of accessible records; (3) accessible records defeat irreversible commitment. The theorem is not definitionally true; it derives recoverability from discrimination.

3.4 Contrapositive and Physical Interpretation

The contrapositive of Theorem 3.2 is immediate and physically decisive.

Corollary 3.3. If irreversible commitment occurs in a physical system, then that system exhibits finite distinguishability.

Since irreversible commitment is a precondition for the existence of facts, records, and empirical observation itself, we conclude that any physically realizable universe must satisfy finite distinguishability.

What This Means: We observe facts every day. You're reading definite words on a page, not a blur of all possible words. Experiments yield specific outcomes. Therefore, by pure logic, our universe must have finite distinguishability. This isn't an assumption—it's a *consequence* of facts existing.

This conclusion is not an empirical hypothesis but a structural necessity. A theory that permits infinite operational refinement cannot support genuine facts; all apparent outcomes remain reversible in principle. Finite distinguishability is therefore not a contingent feature of quantum mechanics but a prerequisite for fact-producing physics of any kind.

3.5 Scope and Consequences

The no-go theorem established here excludes broad classes of candidate physical frameworks, including:

- Infinite-precision classical mechanics
- Purely reversible ontologies in which all distinctions remain physically accessible
- Theories with unbounded operational resolution

The result does not depend on quantum mechanics and applies equally to classical, quantum, or hybrid theories.

---

4. Consequences of Physical Admissibility

In Plain Language: Once we accept that distinguishability is finite and that irreversible commitment happens, several familiar features of physics follow automatically. This section shows that entropy increase (the second law of thermodynamics), the arrow of time, and the fact that measurements have finitely many outcomes aren't separate assumptions—they're logical consequences of having facts at all.

In this section we derive general physical consequences of the admissibility constraints. These results do not assume quantum mechanics or any specific dynamics. They follow solely from finite distinguishability and irreversible commitment.

4.1 Bounded Operational Information Content

Finite distinguishability immediately implies a bound on the amount of operationally accessible information.

Proposition 4.1 (Bounded Information Capacity). Let N be the maximal number of mutually distinguishable states accessible within a bounded system under admissible procedures. Then the operational information content satisfies

$S_{operational} \leq \log N$.

In Plain Language: If you can only tell apart N different states, then you can only store log N bits of information. A coin you can only read as "heads" or "tails" stores 1 bit, no matter how many microscopic configurations might underlie each face.

Proof. By definition, no admissible procedure can reliably distinguish more than N states. Any encoding exceeding this bound necessarily maps multiple input distinctions to the same operational state, resulting in information loss. The maximal Shannon entropy achievable under these constraints is therefore log N. ■

This result does not require discreteness of the mathematical state space. Continuous descriptions remain admissible provided that operational access terminates at finite resolution.

4.2 Entropy Increase as Forced Information Compression

Irreversible commitment corresponds operationally to many-to-one mappings on the admissible state space.

Proposition 4.2 (Entropy Increase Under Commitment). Any admissible process exhibiting irreversible commitment produces non-decreasing operational entropy.

In Plain Language: Every time a fact gets established, some distinctions are erased—multiple possibilities collapse into one outcome. That's information loss from the system's perspective, which means entropy goes up. The second law of thermodynamics isn't a statistical fluke; it's the direct consequence of facts happening.

Proof. Let Φ be an admissible process exhibiting irreversible commitment. Then there exist distinct states $\rho_1 \neq \rho_2$ such that $\Phi(\rho_1) = \Phi(\rho_2)$. Under finite distinguishability, no admissible refinement can recover this distinction. Therefore the number of distinguishable states accessible after Φ is strictly less than before Φ. Since operational entropy is bounded by log N, this reduction corresponds to non-decreasing entropy when viewed as missing information relative to the prior distinguishability structure. ■

Entropy increase is thus not a statistical assumption. It is a structural consequence of finite distinguishability combined with irreversible commitment.

4.3 Emergence of Temporal Ordering

Irreversible commitment induces a partial order on physical processes.

Definition 4.3 (Commitment Ordering). Given two events A and B, we say that A *precedes* B if there exists a sequence of admissible processes from A to B containing at least one irreversible commitment, and no admissible sequence exists from B to A restoring the distinctions committed at A.

Proposition 4.4 (Arrow of Time). The commitment ordering defines a directed temporal structure consistent with the observed arrow of time.

In Plain Language: Why does time have a direction? Why do we remember the past but not the future? Because facts accumulate through irreversible commitment. Once something is recorded, it can't be unrecorded. This creates an asymmetry: the direction of increasing records *is* the direction of time. We don't need to assume time flows forward; it falls out of the logic of fact-making.

Proof. By definition, irreversible commitments cannot be undone within the admissible domain. Therefore the ordering induced by commitment is asymmetric and transitive, yielding a partial order. This ordering coincides with the direction in which records accumulate and entropy increases. ■

This result does not presuppose a fundamental time parameter. Temporal direction emerges from the structure of admissible processes.

Remark on the time parameter in reversible evolution. Section 5 will introduce a parameter t labeling reversible transformations $\{T_t\}$. This might seem circular: if time emerges from commitment, how can we have a pre-existing time parameter? The resolution: the parameter t is merely a label for a one-parameter group of transformations. It does not presuppose temporal direction. The group $\{T_t\}$ could run either way ($t \to -t$ gives another valid parametrization). What commitment provides is the *direction*: the future is the direction in which records accumulate and entropy increases. The parameter t becomes physical time only when oriented by commitment structure.

4.4 Finite Outcome Sets

Operational records correspond to stabilized outcomes of irreversible commitments.

Proposition 4.5 (Finite Outcome Necessity). Any admissible measurement or record-forming process yields outcomes drawn from a finite set.

Proof. Suppose a measurement produced an infinite set of mutually distinguishable outcomes within a bounded system. This would violate finite distinguishability. Therefore admissible record formation necessarily involves finite outcome sets. ■

In Plain Language: When you measure something, you get one of a finite list of possible results—"heads or tails," "red, green, or blue," "spin up or down." You never get infinitely many distinguishable outcomes from a bounded system. This isn't a practical limitation; it's structurally required.

4.5 Exclusion of Infinite-Precision Ontologies

Corollary 4.6. Any physical theory that permits infinite operational refinement is incompatible with irreversible facts, entropy increase, and persistent records.

Such theories may remain mathematically consistent but fail as descriptions of fact-producing physics.

5. Reversible Evolution Between Commitments

In Plain Language: Facts happen at discrete moments—when measurements occur, when records form. But what happens *between* those moments? This section shows that between commitments, physics must evolve in a very specific way: it must preserve all distinctions perfectly (unitary evolution) and be generated by energy (Hamiltonian dynamics). The Schrödinger equation—the master equation of quantum mechanics—isn't assumed; it's *derived* as the only possibility.

In this section we characterize the structure of admissible reversible evolution occurring between irreversible commitments. We establish that admissible reversible evolution must be unitary and generated by a self-adjoint Hamiltonian.

Why must reversible evolution exist? If commitments occurred at every instant—if every moment involved fact-creation—there would be no dynamical evolution at all, only a static sequence of disconnected facts with no law connecting them. But physics succeeds empirically: we can predict future measurement outcomes from present ones. This predictability implies that systems evolve lawfully between commitments. The existence of predictable dynamics is thus evidence that reversible (non-committing) evolution occurs. What follows characterizes the structure such evolution must have.

5.1 Reversible Dynamics as Distinguishability-Preserving Maps

Between irreversible commitments, physical evolution must preserve all operationally accessible distinctions. If reversible dynamics were permitted to erase or amplify distinguishability, then commitment would either occur prematurely or become ill-defined.

In Plain Language: If physics could blur distinctions between measurement events, that blurring would itself constitute a fact—a commitment. But we're talking about periods when no facts are being created. So during these periods, all distinctions must be perfectly preserved. Nothing is gained, nothing is lost—information is just reshuffled.

Let $\{T_t\}_{t \in \mathbb{R}}$ denote a family of admissible reversible transformations acting on the operational state space $\mathcal{S}$. We assume the following minimal properties:

(i) *Composability*: $T_{t+s} = T_t \circ T_s$ for all t, s, with $T_0 = I$.

(ii) *Reversibility*: For each t, T_t admits an admissible inverse T_{-t}.

(iii) *Continuity*: The map $t \mapsto T_t$ is continuous in the operational topology induced by D.

(iv) *Distinguishability preservation*: For all states $\rho, \sigma \in \mathcal{S}$,

$D(T_t(\rho), T_t(\sigma)) = D(\rho, \sigma)$.

These assumptions express the minimal requirement that reversible evolution neither creates nor destroys operational distinctions.

5.2 Deriving the Reversible Sector: From Admissibility to Hilbert Space

The admissibility constraints developed so far—finite operational distinguishability and irreversible commitment—do not by themselves fix a unique mathematical representation for reversible dynamics. They constrain what must not happen (e.g., creation or destruction of operational distinctions), but they leave open multiple possible reversible structures compatible with those constraints.

Meta-theorem (Informal). Admissibility restricts the reversible sector to a narrow family of operational theories. Any operational theory of reversible dynamics that is continuous, compositional, and measurement-complete admits a Jordan algebra representation. Within that family, the postulates R1–R5 below constitute a canonical minimal set that yields complex quantum mechanics. Departures from R1–R5 either (a) collapse into equivalent structure, or (b) introduce operationally inaccessible degrees of freedom, violating the methodological core of admissibility.

We therefore do not claim that R1–R5 are the *unique* choice—only that they form the minimal operationally complete closure of the reversible sector.

5.2.1 Operational Postulates for Reversible Dynamics

We adopt the following postulates, all of which are motivated by admissibility rather than by quantum formalism. Importantly, R1, R4, and R5 are not independent axioms but can be *derived* from minimal assumptions about experimental capabilities; see Appendix E for the full derivations.

R1. Convexity (Operational Mixing). If a preparation procedure can be implemented using classical randomness over admissible preparations, then the resulting state is represented by the corresponding convex mixture.

Derived status: This follows from the existence of classical control registers (Assumption CC, Lemma E.1). Any laboratory capable of recording facts can implement classical randomization over preparations.

R2. Continuous Reversible Dynamics. Between irreversible commitments, admissible dynamics form a continuous, connected group of reversible transformations acting on the operational state space.

This postulate is forced by admissibility. If reversible dynamics were discontinuous, finite changes in experimental control would produce abrupt changes in outcomes, effectively generating new commitments. Continuity is therefore required to prevent reversible evolution from producing facts.

Remark on discrete time. If time is fundamentally discrete (as in some quantum gravity approaches), the argument applies to the discrete unitary step operator U, with the generator becoming a phase operator Φ satisfying $U = e^{-i\Phi}$. The framework is compatible with discrete-time dynamics.

R3. Distinguishability-Preserving Reversibility. All reversible transformations preserve operational distinguishability:

$$D(T(\rho), T(\sigma)) = D(\rho, \sigma).$$

If a reversible transformation erased distinctions, it would constitute an irreversible commitment. If it amplified distinctions, it would generate new operational facts. Reversibility therefore implies isometric action on the distinguishability structure.

R4. Operational Closure of Composition. If two composite states are indistinguishable by all admissible local measurements and their correlations, then they are operationally identical and must be identified in the physical state space.

Equivalently: the joint state of a composite system is completely characterized by the statistics of admissible measurements on its subsystems and their correlations (tomographic locality).

Derived status: This follows from universal composite controllability (Assumption UC, Theorem E.5). If a laboratory can perform all local operations plus a universal entangling gate set, then any state difference not detectable by local measurements and correlations is operationally inaccessible and should be quotiented out.

Remark on hidden variables. Bohmian mechanics posits ontological hidden variables that are operationally inaccessible. This does not conflict with R4: Bohmian particles do not alter operational statistics, so they represent an ontological overlay on quantum structure, not a counterexample to operational closure. The admissibility framework constrains operational physics; ontological interpretations that reproduce quantum statistics remain compatible.

R5. Maximal Reversibility (Purification). Every mixed state arises as the marginal of a pure state of a larger system, and any two purifications of the same state differ by a reversible transformation on the purifying subsystem.

Derived status: This follows from reversible embeddability (Assumption RE, Theorem E.3). The key insight is that stochastic reversible evolution is impossible (Proposition E.2): genuine stochasticity would make the inverse many-to-one, hence irreversible. Therefore, all apparent randomness in mixed states must reflect ignorance about correlations with inaccessible degrees of freedom—which is precisely what purification asserts.

Status. R5 is the weakest derived postulate; relaxing Assumption RE leads to theories with fundamental decoherence, which may be physically relevant in quantum gravity contexts.

Remark on postulate status. R2 and R3 are directly forced by admissibility. R1, R4, and R5 are derived from minimal experimental capability assumptions (CC, UC, RE) as shown in Appendix E.

5.2.2 Consequences of the Postulates

Roadmap. The following sections use the mathematical theory of Jordan algebras to classify all possible reversible sectors compatible with R1–R5. This may seem like an abrupt introduction of abstract mathematics, so we briefly explain the logic:

1. **Why Jordan algebras?** R1–R3 imply the state space is a convex cone with a transitive symmetry group. R4–R5 imply the cone is self-dual (states and effects have matching structure). The Koecher-Vinberg theorem states that such cones are *exactly* the cones of squares in Jordan algebras. This is not an assumption—it's a classification theorem from pure mathematics.
2. **What is a Jordan algebra?** A commutative algebra satisfying $x \circ (y \circ x^2) = (x \circ y) \circ x^2$. The physical content is that observables can be "squared" and "multiplied" in a consistent way. Quantum observables (Hermitian matrices) satisfy this.
3. **Why does classification help?** Jordan, von Neumann, and Wigner proved there are only five types of finite-dimensional Jordan algebras. We can examine each and check whether it satisfies all of R1–R5. Most fail. Complex quantum mechanics is what remains.

This is the mathematical backbone of the uniqueness result. Readers unfamiliar with Jordan algebras can take the classification as given and focus on the exclusion arguments in §5.2.3.

From R1–R3, the operational state space is a finite-dimensional convex set admitting a continuous, transitive group of distinguishability-preserving transformations acting on its extremal (pure) states.

From R4 and R5, the state space is homogeneous and self-dual: states and effects possess matching structure, and composition behaves consistently across subsystems.

By the Koecher–Vinberg theorem [10], any finite-dimensional homogeneous self-dual cone is isomorphic to the cone of squares in a formally real Jordan algebra. The reversible sector must therefore be representable by one of the finite-dimensional Jordan algebras classified by Jordan, von Neumann, and Wigner [11].

The classification yields the following possibilities:

- Real quantum mechanics (Hermitian matrices over $\mathbb{R}$)

- Complex quantum mechanics (Hermitian matrices over ℂ)
- Quaternionic quantum mechanics (Hermitian matrices over ℍ)
- Spin-factor theories
- The exceptional Albert algebra (27-dimensional, over the octonions)

5.2.3 Exclusion of Non-Quantum Foils

The admissibility-motivated postulates exclude all but one of these possibilities.

Exceptional Jordan algebra (Albert algebra). The Albert algebra lacks a consistent tensor-product structure for composing independent systems. This violates tomographic locality (R4) and is therefore inadmissible.

Quaternionic quantum mechanics. Quaternionic theories violate tomographic locality: the joint state of a composite system is not determined by local measurement statistics [3, 24]. They also possess excess global degrees of freedom that are operationally inaccessible, conflicting with admissibility.

Real quantum mechanics. Real quantum mechanics lacks sufficient continuous symmetry to act transitively on pure states (R2 requires connected group action). The transformation group O(n) is disconnected, meaning not all pure states can be continuously connected. RQM also fails to support the full interference structure required by distinguishability-preserving reversible dynamics—relative phases are restricted to ±1 rather than the full circle.

In Plain Language — Why Complex Numbers?

Why does nature use complex numbers (with their mysterious √−1) rather than ordinary real numbers? The answer comes down to *interference* and *continuity*.

In quantum mechanics, when two paths lead to the same outcome, they can interfere—sometimes reinforcing, sometimes canceling. The amount of interference depends on a *phase*, an angle that can take any value from 0° to 360°. With real numbers, you only get two options: +1 (constructive) or −1 (destructive). There's no way to continuously dial between them.

But admissibility requires continuous reversible evolution. If you could only flip discretely between +1 and −1, that flip would itself constitute a fact—a commitment. Continuous phases require the full circle, and the full circle requires complex numbers.

Additionally, real quantum mechanics has a topological problem: some states can't be continuously transformed into others. It's like having a Möbius strip instead of a cylinder—you can't smoothly rotate everything. Complex numbers fix this by providing enough "room" for all transformations to connect continuously.

So complex numbers aren't a mysterious mathematical convenience. They're *forced* by the requirement that reversible evolution be truly continuous and that interference be smoothly adjustable. Real numbers are too rigid; complex numbers are exactly right.

Technical note: The arguments above are heuristic. The rigorous exclusion of real QM relies on *local tomography*: in real quantum mechanics, you can't determine the joint state of two systems from local measurements alone. This violates R4 (operational closure of composition). See Appendix B.4 for the full proof.

Spin-factor theories. These reduce to real or complex quantum mechanics in low dimensions and do not yield distinct admissible theories in the present framework.

The only remaining possibility compatible with R1–R5 is complex quantum mechanics.

5.2.4 Emergence of Hilbert Space as Representation

Lemma 5.1 (Hilbert Space Emergence). Under postulates R1–R5, the operational state space is isomorphic to the space of density operators on a complex Hilbert space $\mathscr{H}$, and reversible transformations are represented by unitary operators acting on $\mathscr{H}$.

In Plain Language: We did not start by assuming particles live in Hilbert space. We started by asking what reversible physics must look like in a universe where facts can exist. Once we require reversible evolution to be continuous, symmetry-respecting, and compatible with how systems compose and how ignorance works, there is only one possible mathematical description left. Hilbert space isn't a guess—it's what remains after everything else is ruled out.

Proof. By the analysis of §5.2.2–5.2.3, the state space must be isomorphic to density matrices on $\mathbb{C}^N$ for some finite N (determined by the distinguishability bound). Reversible transformations preserving distinguishability and satisfying R2–R3 must act as unitary conjugations: $T(\rho) = U\rho U^\dagger$ for unitary U. See Appendix B.4 for the complete proof. ■

5.2.5 Scope Clarification

It is important to emphasize what has and has not been shown.

What admissibility alone forces: The constraints of finite distinguishability and irreversible commitment do not by themselves select a unique reversible structure. They do, however, motivate and constrain the operational postulates R1–R5.

What the postulates force: Under R1–R5, complex Hilbert space quantum mechanics is uniquely selected. Other structures (real, quaternionic, exceptional) are excluded by specific postulates.

The logical structure: Admissibility → motivates R1–R5 → forces Hilbert space.

Hilbert space is thus not fundamental ontology, but the mathematically faithful representation of admissible reversible dynamics in a fact-producing universe.

5.3 Unitary Implementation of Reversible Evolution

Theorem 5.1 (Unitary Realization). Let $\{T_t\}$ be a continuous, reversible, distinguishability-preserving family of admissible transformations on a state space satisfying the conditions of Lemma 5.1. Then there exists a complex Hilbert space $\mathscr{H}$ and a strongly continuous one-parameter unitary group $\{U(t)\}$ acting on $\mathscr{H}$ such that

$$T_t(\rho) = U(t)\, \rho\, U(t)^\dagger.$$

In Plain Language: "Unitary" means the transformation preserves all lengths and angles—nothing is stretched, squashed, or lost. We've just proven that between measurements, physics *must* be unitary. This is usually taken as an axiom of quantum mechanics. Here, it's a theorem.

Proof. By Lemma 5.1, the state space admits a representation as density matrices on a Hilbert space $\mathscr{H}$. Distinguishability preservation implies that each T_t preserves the trace distance:

$$\|T_t(\rho) - T_t(\sigma)\|_1 = \|\rho - \sigma\|_1.$$

By Wigner's theorem [12], any bijection on pure states preserving transition probabilities (equivalently, trace distance for pure states) is implemented by either a unitary or antiunitary operator.

Since $\{T_t\}$ forms a continuous one-parameter group and $T_0 = I$, the implementing operators must vary continuously with t. Antiunitary operators are disconnected from the identity (they reverse orientation on the Hilbert space), so continuity forces each T_t to be implemented by a unitary operator U(t).

The group property $T_{t+s} = T_t \circ T_s$ implies $U(t+s) = e^{i\varphi(t,s)}\, U(t)U(s)$ for some phase φ. By standard arguments [13], continuity allows the phase to be absorbed into the definition of U(t), yielding a strongly continuous unitary representation. ■

5.4 Existence of the Hamiltonian Generator

Theorem 5.2 (Hamiltonian Necessity). Any strongly continuous one-parameter unitary group $\{U(t)\}$ admits a unique self-adjoint generator H such that

$$U(t) = \exp(-iHt/\hbar).$$

In Plain Language: If evolution is unitary and continuous, there must be something generating it—and that something is what we call energy (the Hamiltonian). This is Stone's theorem from mathematics. The Hamiltonian isn't an extra assumption; it exists automatically once you have continuous unitary evolution.

Proof. This is Stone's theorem [14]. Strong continuity of $t \mapsto U(t)$ guarantees that the limit

$H\psi = i\hbar \lim_{t\to 0} [U(t)\psi - \psi]/t$

exists on a dense domain and defines a self-adjoint operator. The generator H is unique up to additive constants corresponding to global phase. ■

The operator H is identified as the Hamiltonian. Its existence is not an independent axiom but a representation-theoretic necessity arising from admissibility.

5.5 Schrödinger Evolution as Admissible Dynamics

Differentiating the unitary evolution yields

$i\hbar\, d\rho/dt = [H, \rho]$

for mixed states, and

$i\hbar\, d|\psi\rangle/dt = H|\psi\rangle$

for pure states. These equations are not postulates but identities expressing the infinitesimal action of admissible reversible evolution.

In Plain Language: The Schrödinger equation—the most famous equation in quantum mechanics—appears here not as an assumption but as the *only possible* way to describe continuous reversible evolution that preserves distinguishability. It's mathematically forced.

The Schrödinger equation thus appears as the unique admissible continuous-time description of reversible evolution between commitments.

---

6. Measurement as Irreversible Commitment

In Plain Language: Measurement has always been the most controversial part of quantum mechanics. What makes it special? Why does it seem to "collapse" the wave function? This section shows that measurement isn't special at all—it's just the moment when irreversible commitment happens. The mathematical structures physicists use to describe measurement (CPTP maps, POVMs, state update rules) aren't arbitrary; they're the minimal description of what commitment does.

In this section we demonstrate that measurement does not require any special dynamical postulate beyond irreversible commitment. The standard mathematical structures of quantum measurement—completely positive maps, POVMs, and state-update rules—are derived as necessities.

6.1 Measurements as Admissible Processes

A measurement is an admissible process whose defining feature is the production of a stable record. Such a record encodes the outcome in a form that cannot be erased by any admissible procedure. Measurement therefore constitutes an instance of irreversible commitment.

Formally, a measurement is represented by an admissible map Φ acting on the operational state space such that distinct pre-measurement states may map to the same post-measurement state while yielding a recorded outcome label i belonging to a finite outcome set.

6.2 Completely Positive Maps and Physical Admissibility

Admissible physical processes acting on quantum states must preserve positivity and normalization under all possible extensions.

Theorem 6.1 (CPTP Necessity). Any admissible irreversible process acting on quantum states is represented by a completely positive, trace-preserving (CPTP) map.

In Plain Language: "Completely positive" is a technical condition that ensures probabilities remain sensible (non-negative) even when your system is entangled with something else. "Trace-preserving" means total probability stays equal to 1. These aren't extra assumptions—they're requirements for any process that could actually happen physically.

Proof. Let Φ be an admissible process on system A. Consider an arbitrary extension where A is embedded in a larger system AB, with B serving as an ancilla that undergoes no direct evolution.

(i) *Positivity*: Φ must map positive operators (valid states) to positive operators. Otherwise, $\Phi(\rho)$ could have negative eigenvalues, which are unphysical.

(ii) *Complete positivity*: The extended map $\Phi \otimes I_B$ must also preserve positivity. If Φ were merely positive but not completely positive, there would exist entangled states ρ_{AB} such that $(\Phi \otimes I_B)(\rho_{AB})$ has negative eigenvalues [15]. Such outputs are not valid quantum states, violating admissibility.

(iii) *Trace preservation*: $\mathrm{Tr}(\Phi(\rho)) = \mathrm{Tr}(\rho) = 1$ is required for normalization of probability.

These conditions are necessary for physical realizability. ■

6.3 Kraus Representation and Outcome Structure

By the Kraus representation theorem [16], any CPTP map Φ admits a decomposition

$$\Phi(\rho) = \Sigma_k M_k \rho M_k^\dagger,$$

where the Kraus operators $\{M_k\}$ satisfy $\Sigma_k M_k^\dagger M_k = I$.

When the process produces a recorded outcome i, the Kraus operators can be grouped into subsets corresponding to distinct outcomes. The associated positive operators

$E_i = \Sigma_{k \in i} M_k^\dagger M_k$

form a positive operator-valued measure (POVM), satisfying $\Sigma_i E_i = I$ and $E_i \geq 0$.

Proposition 6.2 (Finite Outcome Necessity). The outcome set of any admissible measurement is finite.

Proof. By finite distinguishability, no admissible procedure can reliably distinguish an infinite set of mutually exclusive outcomes within a bounded system. Therefore the POVM must consist of finitely many effects. ■

6.4 State Update as Minimal Commitment

Conditioning on a specific measurement outcome i yields the post-measurement state

$\rho_i = (M_i \rho M_i^\dagger) / \mathrm{Tr}(E_i \rho)$,

where M_i denotes the effective Kraus operator associated with outcome i.

Theorem 6.3 (No-Extra-Disturbance Principle). The state update rule contains no additional dynamical disturbance beyond irreversible commitment.

In Plain Language: The infamous "wave function collapse" isn't a mysterious additional process. It's just what happens when you update your description after learning which outcome occurred. The math of commitment already contains the state update rule. Nothing extra is needed.

Proof. The update rule arises purely from conditioning on the observed outcome within the CPTP structure. No supplementary stochastic or dynamical modification is required. All loss of distinguishability is accounted for by the many-to-one nature of commitment itself. ■

This result shows that the projection postulate can be understood as a bookkeeping rule encoding commitment rather than an independent physical process. The apparent "collapse" reflects the transition from the pre-commitment ensemble to the post-commitment conditioned state.

6.5 POVMs as Finite-Resolution Necessity

Idealized projective measurements correspond to the special case where POVM elements are orthogonal projectors. In realistic physical situations, finite resolution generically yields non-projective POVMs.

Proposition 6.4 (Generic POVM Structure). Finite distinguishability implies that admissible measurements are generically represented by POVMs rather than sharp projective measurements.

Proof. If operational resolution is finite, measurement effects necessarily coarse-grain over underlying degrees of freedom. This coarse-graining produces positive operators that are not idempotent, yielding POVM structure. ■

---

7. Probability and the Born Rule

In Plain Language: Where do quantum probabilities come from? Why do we square the wave function to get probabilities (the Born rule)? This section proves that the Born rule isn't an arbitrary choice—it's the *only* way to assign probabilities that's compatible with unitary evolution and finite distinguishability. Any other rule would lead to contradictions.

In this section we derive the Born rule as the unique probability assignment compatible with admissibility constraints.

7.1 The Structural Role of Probability

Probability enters physics only at the interface between reversible evolution and irreversible commitment. Prior to commitment, admissible dynamics preserves all operational distinctions and no outcome is selected. At commitment, multiple possibilities are reduced to a single recorded fact. Probability quantifies the relative weights with which outcomes are realized.

In Plain Language: Probabilities don't exist during the quantum evolution—all possibilities are still "in play." Probability only becomes meaningful when a fact is established, at the moment of commitment. Then it tells you how likely each outcome was.

Crucially, probability is not interpreted here as subjective belief. It is an objective feature of admissible outcome statistics arising from the structure of reversible dynamics constrained by finite distinguishability.

7.2 Constraints on Admissible Probability Assignments

Let $\{E_i\}$ be the finite POVM associated with an admissible measurement. An admissible probability assignment $p(i|\rho)$ must satisfy:

(i) *Positivity*: $p(i|\rho) \geq 0$ for all i.

(ii) *Normalization*: $\Sigma_i\, p(i|\rho) = 1$.

(iii) *Additivity under coarse-graining*: Probabilities assigned to coarse-grained outcomes equal the sum of probabilities of the fine-grained outcomes they comprise.

(iv) *Unitary invariance*: $p(i|U\rho U^\dagger) = p(U^\dagger E_i U|\rho)$ for all unitaries U.

(v) *Continuity*: Small changes in ρ induce small changes in $p(i|\rho)$.

(vi) *Non-contextuality*: The probability for outcome i depends only on ρ and E_i, not on other elements of the POVM or the physical implementation.

7.3 Uniqueness of Quadratic Probability

Theorem 7.1 (Born Rule Uniqueness). Let ρ be a quantum state and $\{E_i\}$ a POVM on a Hilbert space of dimension $d \geq 3$. Then the unique probability assignment satisfying the admissibility constraints is

$$p(i|\rho) = \mathrm{Tr}(E_i\, \rho).$$

In Plain Language: This is the Born rule: probability equals trace of (effect times state), which for pure states gives you the squared amplitude. The theorem says this is the *only* formula that works. Anything else would violate one of the basic requirements like additivity or non-contextuality.

Proof. We establish this result in three steps.

Step 1: Frame functions and Gleason's theorem. For projective measurements in dimension $d \geq 3$, Gleason's theorem [17] states that any non-contextual probability assignment must take the form $p(P|\rho) = \mathrm{Tr}(P\rho)$ for projectors P. The key insight is that non-contextuality and additivity over orthogonal projectors leave no freedom: any deviation from the trace rule produces contradictions when measurements are refined or coarse-grained.

Step 2: Extension to POVMs. Busch [18] and Caves et al. [19] extended Gleason's theorem to POVMs. Any additive, non-contextual probability measure on the space of effects (positive operators E with $0 \leq E \leq I$) must be linear in the state: $p(E|\rho) = \mathrm{Tr}(E\rho)$.

Step 3: Admissibility forces non-contextuality. This step requires careful argument, as the relationship between operational equivalence and non-contextuality has subtleties [20].

Consider a POVM element E_i that appears in two different POVMs: $\{E_1, E_2, ..., E_n\}$ and $\{E_1, F_2, ..., F_m\}$. A contextual probability assignment would assign different values $p(E_1|\rho)$ depending on whether E_1 is measured alongside $\{E_2, ...\}$ or alongside $\{F_2, ...\}$.

The reproducibility argument. Suppose probabilities were contextual: $p(E_1|\rho, \mathrm{POVM}_1) \neq p(E_1|\rho, \mathrm{POVM}_2)$. Then over many trials, the observed frequencies would differ systematically between the two measurement contexts. But this frequency difference is itself an empirical regularity—a *stable, reproducible operational fact*.

By the definition of operational fact (§2), any such regularity must be traceable to some physically accessible distinction. Either:

(a) The distinction is measurable, in which case M and M' are operationally distinguishable (contrary to the assumption that they share E_i), or

(b) The distinction is unmeasurable, in which case it cannot ground reproducible statistics without violating admissibility.

Option (b) would require an inaccessible degree of freedom to reliably influence outcome frequencies across repeated trials. But "reliably influence across repeated trials" means producing a stable operational pattern—exactly what admissibility says inaccessible degrees of freedom cannot do.

Formal statement. Let M and M' be two measurement procedures that both include effect E_i. Suppose $p(E_i|\rho, M) \neq p(E_i|\rho, M')$. By finite distinguishability, the difference must be encoded in some accessible degree of freedom. But the outcome records for "E_i clicked" are operationally identical in both cases—they produce the same persistent fact. Therefore any degree of freedom that distinguishes M from M' is inaccessible to the outcome record, and by admissibility, cannot influence the probability of that outcome.

This argument shows that *operational* non-contextuality—probabilities depending only on the effect and state, not on co-measured effects—is forced by admissibility. Note that this does not rule out *ontological* contextuality in the sense of Spekkens [20], where underlying ontic states might differ; it rules out *observable* contextuality in outcome statistics.

Alternative derivation (Appendix E). A cleaner version of this argument defines an "operational effect" as an equivalence class of outcome-events with identical statistics on all states. Under this definition, non-contextuality is automatic: if probabilities depended on context, the outcome-events would belong to different equivalence classes and would not be "the same effect" (Lemma E.6).

Therefore non-contextuality is forced, and by the generalized Gleason theorem, $p(i|\rho) = \mathrm{Tr}(E_i \rho)$. ■

Remark on dimension. For d = 2 (qubits), Gleason's theorem fails, and additional probability assignments exist. However, these alternatives violate either continuity or compositionality when qubits are embedded in larger systems [20]. Since admissibility requires consistent probability assignments across all system sizes, the Born rule is unique even in the qubit case.

7.4 Operational Interpretation

The quadratic form of the Born rule reflects the preservation of distinguishability under unitary evolution. Squared amplitudes correspond to invariant measures on the space of admissible states. Alternative assignments would permit amplification or suppression of operational distinctions under reversible evolution, contradicting admissibility.

From this perspective, probability measures the fraction of admissible distinguishability capacity allocated to each outcome channel at commitment.

---

8. Entanglement and Nonlocal Correlations

In Plain Language: Entanglement is often called the strangest feature of quantum mechanics—"spooky action at a distance," as Einstein put it. Two particles can be correlated in ways that seem impossible if each particle has its own independent reality. This section shows that entanglement isn't a weird add-on; it's *inevitable* once you have quantum structure. And crucially, it doesn't involve any faster-than-light influence. The "spookiness" is in how facts get established globally, not in any signal being sent.

In this section we establish that entanglement and Bell correlations arise as structural consequences of admissibility, without requiring violations of locality.

8.1 Composite Systems and Tensor Structure

Consider two subsystems A and B, each described by an admissible operational state space. The joint system is described by a composite state space whose reversible dynamics factorize as tensor products of admissible unitary evolutions:

$$U_{AB}(t) = U_A(t) \otimes U_B(t)$$

for independent evolutions. This compositional structure is forced by admissibility: reversible dynamics on independent systems must compose without creating or destroying distinguishability.

Entangled states are those joint states that cannot be written as convex combinations of product states:

$$\rho_{AB} \neq \Sigma_k p_k \rho_A^{(k)} \otimes \rho_B^{(k)}.$$

8.2 Shared Distinguishability Structure

In Plain Language: Think of two coins that are secretly connected. Each coin individually looks random—50% heads, 50% tails. But together, they're always correlated: if one is heads, the other is tails. The information about their relationship isn't stored in either coin alone; it's stored in their relationship. That's entanglement. The "information" about correlations is nonlocal—distributed across the pair—but no signal travels between them.

In the admissibility framework, distinguishability is a relational resource. For composite systems, distinguishability need not decompose additively. Entangled states correspond to configurations where distinguishability is stored nonlocally across the composite state space.

Operationally, while local reduced states may be indistinguishable from classical mixtures, global measurements can access distinctions that no local procedure can resolve independently.

8.3 Bell Correlations Without Dynamical Nonlocality

Bell inequalities [21] constrain correlations achievable under local hidden-variable models where outcomes are locally determined prior to measurement.

Theorem 8.1 (Bell Compatibility). Admissible theories satisfying finite distinguishability, unitary reversible dynamics, and irreversible commitment can violate Bell inequalities without permitting superluminal signaling.

In Plain Language: Bell's theorem shows that quantum correlations are stronger than any "local realistic" theory could produce—there's no way to explain them by saying each particle secretly carried its answer all along. But violating Bell inequalities doesn't require faster-than-light communication. The trick is that outcomes aren't determined until commitment happens, and commitment is global. No signal passes; it's the *structure of fact-creation* that's nonlocal.

Proof. Prior to commitment, entangled states encode joint distinguishability structure without definite local outcomes. The key observation is that outcome determination occurs at commitment, which is global with respect to the composite state.

Consider a Bell experiment with entangled state $|\psi\rangle_{AB}$ and local measurement settings a for Alice and b for Bob. The joint probability is

$$p(i,j|a,b) = \mathrm{Tr}[(E_i^a \otimes E_j^b)|\psi\rangle\langle\psi|].$$

This probability is determined by the global state and the tensor product of local effects. Since commitment selects correlated outcomes according to the Born rule, and since the entangled state encodes non-factorizable correlations, Bell inequalities can be violated.

However, the marginal probabilities satisfy

$$p(i|a) = \Sigma_j\, p(i,j|a,b) = \mathrm{Tr}[(E_i^a \otimes I)|\psi\rangle\langle\psi|] = \mathrm{Tr}[E_i^a\, \rho_A],$$

which is independent of Bob's setting b. No superluminal signaling is possible. ■

8.4 No-Signaling as an Admissibility Constraint

Proposition 8.2 (No-Signaling). For any bipartite state and any local admissible operation on subsystem A, the marginal outcome probabilities on subsystem B remain invariant.

Proof. Let Φ_A be a local admissible operation (CPTP map) on A. The joint state transforms as $\rho_{AB} \mapsto (\Phi_A \otimes I_B)(\rho_{AB})$. The reduced state on B is

$$\rho_B' = Tr_A[(\Phi_A \otimes I_B)(\rho_{AB})] = Tr_A[\rho_{AB}] = \rho_B,$$

where the last equality follows from the trace-preserving property of Φ_A. Since measurement statistics on B depend only on ρ_B, they are unaffected by operations on A. ■

8.5 Structural Interpretation

Entanglement reflects that admissible outcome commitment applies to the composite system as a whole. Correlations are fixed at the level of global commitment, not propagated dynamically between subsystems. This interpretation preserves locality of admissible dynamics while explaining the observed nonlocal correlations.

---

9. Extension to Relativistic Admissibility

In Plain Language: So far, we've derived non-relativistic quantum mechanics. But what about Einstein's relativity? This section shows that when you require physics to respect finite signal speed (nothing faster than light) and work the same way everywhere, you're forced into a very specific structure. Particles must have "spin"—an intrinsic angular momentum—and obey the Dirac equation. Again, this isn't assumed; it follows from admissibility plus locality.

Important clarification on scope. The results of Sections 2–8 derive quantum structure from admissibility constraints alone, without invoking relativity. This section extends the framework by adding a new physical requirement: *relativistic locality*, meaning that admissible dynamics must be compatible with finite signal propagation and Lorentz invariance. This is an additional empirical input, not a consequence of admissibility per se. The goal is to show that once relativistic locality is imposed, the Dirac equation emerges as uniquely compatible with admissible quantum dynamics—not to derive relativity from admissibility.

In this section we demonstrate that requiring first-order local dynamics forces Clifford algebra structure and the Dirac equation.

9.1 First-Order Local Evolution

Non-relativistic Schrödinger evolution is second-order in spatial derivatives. A relativistic theory must admit first-order local evolution compatible with finite signal propagation.

We impose two additional constraints beyond the core admissibility framework:

(i) *First-order locality*: Admissible reversible evolution is generated by operators linear in the generators of spatial translation (momenta). This ensures that the evolution equation treats space and time symmetrically (both first-order), as required by Lorentz covariance.

(ii) *Square consistency*: The square of the first-order generator reproduces the relativistic dispersion relation $E^2 = p^2 + m^2$. This encodes the relativistic relationship between energy and momentum.

In Plain Language: Relativity requires that the laws of physics involve space and time on equal footing—first derivatives in time should match first derivatives in space. The Schrödinger equation is second-order in space but first-order in time, so it's not truly relativistic. We need a first-order-in-everything equation. But there's a constraint: the square of this first-order equation should give back the correct relationship between energy and momentum.

These constraints are motivated by special relativity, which is treated here as an empirical fact about our universe rather than derived from admissibility. The question this section addresses is: *given* that relativity holds, what does admissibility require of relativistic quantum dynamics?

9.2 Algebraic Forcing of Clifford Structure

Let D denote a first-order generator. Square consistency requires

$$D^2 = -\Delta + m^2 = p^2 + m^2,$$

where Δ is the Laplacian.

Theorem 9.1 (Clifford Forcing). Any first-order operator D whose square yields a scalar second-order invariant must be representable in terms of matrices satisfying a Clifford algebra.

Proof. Write $D = \Sigma_i \alpha_i p_i + \beta m$ for matrices α_i, β. Computing D^2:

$$D^2 = \Sigma_{ij} \alpha_i \alpha_j p_i p_j + m \Sigma_i (\alpha_i \beta + \beta \alpha_i) p_i + \beta^2 m^2.$$

For this to equal $p^2 + m^2 = \Sigma_i p_i^2 + m^2$, we require:

$$\alpha_i \alpha_j + \alpha_j \alpha_i = 2\delta_{ij} I, \alpha_i \beta + \beta \alpha_i = 0, \beta^2 = I.$$

These are precisely the defining relations of a Clifford algebra $Cl_{1,3}(\mathbb{R})$. No alternative algebraic structure satisfies square consistency while preserving linearity. ■

In Plain Language: To make a first-order equation that squares to the right thing, you need matrices that anticommute in a specific pattern. This pattern is called a Clifford algebra. There's no other option—mathematics forces your hand.

9.3 Emergence of Spinors

Corollary 9.2 (Spin Necessity). Any admissible first-order relativistic quantum theory must involve spinorial degrees of freedom.

Proof. The minimal faithful representation of the Clifford algebra $Cl_{1,3}(\mathbb{R})$ is four-dimensional, realized by the Dirac matrices γ^u. States transforming under this representation are four-component spinors. Scalar or vector representations do not support first-order square-consistent dynamics. ■

In Plain Language: Spinors are mathematical objects that describe particles with spin, like electrons. The theorem says you *can't avoid* spinors if you want relativistic quantum mechanics. Spin isn't an empirical surprise tacked on to the theory—it's required by the mathematics of first-order relativistic evolution.

9.4 Dirac Equation as Admissible Dynamics

With Clifford structure fixed, the unique admissible first-order evolution equation is

$(i\gamma^u \partial\_\mu - m)\psi = 0.$

This is the Dirac equation. Lorentz covariance emerges as a consequence of the algebraic structure rather than an independent postulate.

In Plain Language: The Dirac equation—which correctly describes electrons and predicted antimatter—is the *only* equation you can write down that satisfies our requirements. Dirac discovered it in 1928 by inspired guesswork. Here, we derive it as the unique possibility.

9.5 Dimensional Considerations

The derivation assumes three spatial dimensions. This assumption is empirically motivated and consistent with the observed universe. Whether admissibility constrains dimensionality is an open question. In d spatial dimensions, Clifford algebras $Cl_{1,a}(\mathbb{R})$ have representation dimensions that depend on d mod 8 (Bott periodicity). The four-dimensional Dirac representation is specific to d = 3. Extension to arbitrary dimension would require additional principles beyond those developed here.

---

10. Quantum Mechanics as an Admissibility Fixed Point

In Plain Language: This section brings everything together. We've shown, step by step, that each piece of quantum mechanics follows from the requirement that facts can exist. There's no wiggle room: start from fact-producing physics, and you end up at quantum mechanics. It's a "fixed point"—the unique theory that satisfies all the constraints simultaneously.

We now synthesize the results and state the central claim precisely.

10.1 Synthesis of Results

The argument proceeds in strictly layered fashion:

1. **Facts require commitment**: Physical facts require irreversible commitment (Section 2.4).
2. **Commitment requires finite distinguishability**: Theorem 3.2 proves that infinite distinguishability is incompatible with irreversible commitment.
3. **Reversible evolution is unitary**: Between commitments, admissible dynamics must preserve distinguishability, forcing unitary evolution (Theorem 5.1).
4. **Hamiltonians exist necessarily**: Stone's theorem yields self-adjoint generators (Theorem 5.2).
5. **Measurement is minimal commitment**: CPTP maps, POVMs, and state update emerge from admissibility (Section 6).
6. **Born rule is unique**: Gleason-type arguments fix the probability rule (Theorem 7.1).
7. **Entanglement follows from global commitment**: Bell correlations arise without nonlocality (Section 8).
8. **Relativistic structure is forced**: First-order locality yields Dirac dynamics (Section 9).

10.2 Fixed-Point Characterization

Theorem 10.1 (Admissibility Fixed Point). Quantum mechanics is the unique physically admissible framework satisfying:

(i) Finite operational distinguishability

(ii) Irreversible commitment producing stable facts

(iii) Reversible dynamics preserving distinguishability

(iv) Compositional locality

(v) Continuous implementability

Any theory violating one or more conditions fails to support fact-producing physics. Any theory satisfying all conditions is unitarily equivalent to quantum mechanics in its reversible sector and exhibits the same measurement and probability structure.

In Plain Language: This is the main theorem. It says: if you want a universe with facts, you get quantum mechanics. Not approximately, not as one option among many, but uniquely and necessarily. Every alternative either can't produce facts or turns out to be quantum mechanics in disguise.

10.3 What Is Explained and What Is Not

Explained:

- Unitary reversible dynamics
- Hamiltonian time evolution
- Measurement as irreversible commitment
- Finite outcome sets
- Quadratic probability (Born rule)
- Entanglement and nonlocal correlations
- Spinorial relativistic structure

Not explained:

- Specific interaction Hamiltonians
- Coupling constants
- Particle spectra
- Numerical values of physical constants
- Gravitational dynamics

In Plain Language: We've explained *why* quantum mechanics has the structure it does. We haven't explained *which* quantum theory describes our universe—that requires additional input about what particles exist, how they interact, and what the constants of nature are. Admissibility tells you the rules of the game; it doesn't tell you who's playing or what the score is.

These features require additional physical input beyond admissibility.

10.4 Significance

The significance of this result is conceptual rather than predictive. It relocates the foundations of quantum mechanics from postulates to necessity. Quantum mechanics no longer appears as a mysterious departure from classical reasoning, but as the inevitable structure of any universe capable of producing irreversible facts.

In Plain Language: Quantum mechanics has seemed strange for a hundred years because we thought it was one possible theory that happened to be true. This paper suggests it's not contingent at all—it's the *only* way physics can work if facts are to exist. The strangeness isn't a bug; it's a feature required by the very existence of a factual world.

---

11. Scope, Limitations, and Open Directions

11.1 Scope of Results

The results are structural and constraint-based. They identify necessary conditions for any physically admissible theory. Within this scope, the mathematical architecture of quantum mechanics is shown to be non-optional.

These results apply independently of interpretation. They do not depend on assumptions about observers, consciousness, branching worlds, or hidden variables.

11.2 Limitations

Admissibility alone does not determine:

- Specific interaction Hamiltonians
- Coupling constants and particle spectra
- The number of fermion generations
- Gravitational dynamics
- Spatial dimensionality (assumed to be 3)

In Plain Language: This framework explains the *form* of physical law, not its *content*. Why there are three generations of quarks, why gravity is so weak, why we live in three spatial dimensions—these remain open questions requiring additional principles.

Remark on infinite-dimensional Hilbert spaces. The Jordan classification and Lemma 5.1 apply to finite-dimensional systems. Actual quantum mechanics routinely uses infinite-dimensional Hilbert spaces (position, momentum, field modes). How does the framework extend?

The resolution lies in the operational interpretation of finite distinguishability. In any bounded region with finite energy, the Bekenstein bound limits the number of distinguishable states to a finite N. Infinite-dimensional Hilbert space emerges only in idealized limits—unbounded regions, infinite energy, or mathematical convenience. The operational distinguishability relevant to fact-production is always finite; infinite dimensions are a mathematical representation that extends smoothly from the finite case.

Alternatively: the framework applies directly to bounded subsystems; extension to unbounded or field-theoretic systems requires additional analysis of how finite distinguishability emerges from holographic or thermodynamic constraints. This connection is discussed further in Appendix D.2.

11.3 Relation to Quantum Field Theory

The admissibility framework is compatible with quantum field theory. The kinematic and measurement-theoretic results apply to quantum fields. Extension to interacting theories requires careful treatment of locality, renormalization, and continuum limits.

11.4 Relation to Gravity

The present work does not propose a theory of gravity. However, any viable theory of gravity must respect admissibility constraints.

11.5 Relation to Everettian Interpretations

The admissibility framework treats fact-existence as an operational precondition for empirical physics. But certain interpretations of quantum mechanics—particularly Everettian (many-worlds) approaches—deny that facts in our sense exist *fundamentally*. In these interpretations, what appear to be irreversible facts are branch-relative phenomena: the wavefunction never collapses, all outcomes occur on different branches, and "irreversibility" reflects the practical inaccessibility of other branches rather than genuine ontological commitment.

How does the admissibility framework relate to Everettian physics?

The key observation is that our constraints apply at the *operational* level, not the ontological level. Even in an Everettian universe, observers within a branch see definite outcomes, form records, and accumulate entropy. From the perspective of any observer capable of doing physics, admissibility constraints hold: distinguishability is finite, commitments are irreversible *relative to the accessible branch*, and probabilities follow the Born rule.

The framework therefore does not adjudicate between interpretations. It shows that *any* interpretation—including Everettian ones—must respect the operational structure of quantum mechanics. If Everettian physics is correct, then "irreversible commitment" is branch-relative rather than absolute. But the *structural* requirements (finite distinguishability, CPTP dynamics, Born rule) remain the same. The admissibility framework identifies what must be true operationally; it leaves open what this means ontologically.

11.6 Relation to QBism

QBism interprets the quantum state as an agent's personal degrees of belief about the consequences of their actions on the world, rather than as an objective physical field. The admissibility framework is compatible with this stance, because it does not require a particular ontology for the quantum state.

What admissibility constrains is the *structure* any agent-facing physical theory must possess in a world where stable records form and irreversible commitments occur. In this sense, admissibility operates one layer "below" interpretational semantics: whether probabilities are read as objective propensities or subjective credences, coherence across reversible evolution and irreversible record formation imposes strong constraints on allowable probability assignments and state-update rules.

The main results (unitarity between commitments, CPTP measurement structure, and Born-rule form) can thus be read as constraints on the calculus any agent must use to remain consistent with the operational architecture of record formation—consistent with QBism's emphasis on normative structure, while remaining neutral on what the quantum state *is*.

Put differently: QBism is an interpretation of the quantum formalism; admissibility is an argument for why the formalism's core architecture is non-optional for any empirically usable physics. The two are orthogonal: QBism addresses what probabilities *mean*, while admissibility addresses why probabilities must take the *form* they do.

11.7 Relation to Classical Statistical Mechanics

A potential objection: classical statistical mechanics produces irreversible facts (thermodynamic equilibration, definite measurement outcomes) without invoking quantum structure. Does this show that classical physics can satisfy admissibility constraints?

The response is that classical statistical mechanics achieves irreversibility through *coarse-graining*, which implicitly assumes finite distinguishability at the operational level. When we describe a gas by its temperature rather than the positions and momenta of 10^{23} particles, we are acknowledging that those microscopic degrees of freedom are operationally inaccessible. The "irreversibility" of thermodynamic processes is exactly the pattern analyzed in Section 3: many microscopic configurations map to the same macroscopic record.

A skeptic might argue that classical chaos can amplify microscopic differences into macroscopic distinctions, allowing facts to form without quantum mechanics. But this argument actually supports the admissibility framework: chaotic amplification creates facts precisely by converting continuous degrees of freedom into discrete, coarse-grained records. The chaotic system functions as a natural analog-to-digital converter—and the moment of digitization is the moment of commitment.

The question is whether classical mechanics at the fundamental level—without coarse-graining—can support facts. Theorem 3.2 argues not: if microscopic classical states are perfectly distinguishable, then any "commitment" is in principle reversible by accessing those microscopic degrees of freedom. Classical statistical mechanics works because it operates in a regime where quantum effects enforce finite distinguishability at small scales, even if the effective description is classical. The admissibility framework explains *why* coarse-graining leads to irreversibility: it is not a failure to track details, but a structural feature of any fact-producing physics.

11.8 The Selection Problem

The framework developed here explains the *structure* of quantum measurement: why outcomes belong to finite sets, why probabilities follow the Born rule, why measurement involves irreversible commitment. But it does not explain *which* outcome is selected from among the possibilities.

This is important to acknowledge. The "measurement problem" has two components: (1) why does measurement have the structural features it does? and (2) why does a specific outcome occur rather than another? The admissibility framework addresses (1) comprehensively. It leaves (2) as a separate question.

On the selection question, the framework is agnostic. It is compatible with interpretations where selection is:

- Ontologically random (Copenhagen-style collapse)
- Deterministic but epistemically inaccessible (hidden variables)
- Only apparently singular, with all outcomes realized in different branches (Everett)
- A primitive feature of nature requiring no further explanation

What the framework establishes is that *whatever* the selection mechanism, it must operate within the constraints of admissibility: finite outcome sets, Born-rule statistics, CPTP evolution. The selection mechanism cannot be contextual, cannot violate no-signaling, and cannot depend on operationally inaccessible degrees of freedom (else it would conflict with finite distinguishability).

In this sense, the framework does not "dissolve" the measurement problem entirely. It reframes it: the structural aspects of measurement are necessary features of fact-producing physics, while the selection aspect remains an open interpretive and potentially empirical question.

11.9 Open Directions

- Investigating whether admissibility constrains spatial dimensionality
- Deriving interaction structures consistent with admissibility
- Exploring connections to holographic entropy bounds
- Examining emergence of classical spacetime from admissible quantum dynamics
- Clarifying the role of admissibility in cosmological initial conditions

11.10 Falsifiability and Testable Implications

A natural concern: if the framework explains existing physics without predicting new phenomena, is it falsifiable?

Status of the framework. The admissibility framework is a *constraint theorem*, not a predictive model. It identifies necessary conditions for fact-producing physics. As such, it is falsifiable only indirectly: any observed violation of its structural predictions would refute the framework.

Structural predictions that could be falsified:

1. *Finite distinguishability bounds.* If experiments could reliably distinguish arbitrarily many states within bounded systems (violating Bekenstein-type bounds), the framework would fail. Current physics shows no such violations; holographic bounds suggest fundamental limits.
2. *Born rule statistics.* If outcome frequencies systematically deviated from $\mathrm{Tr}(E\rho)$ in ways not attributable to experimental error, the uniqueness proof of Theorem 7.1 would be refuted. A century of quantum experiments confirms Born statistics.
3. *CPTP dynamics.* If measurement processes violated complete positivity or trace preservation in reproducible ways, the measurement derivation of Section 6 would fail.
4. *No-signaling.* If entangled systems permitted faster-than-light communication, Proposition 8.2 would be falsified.

Operationally testable implications:

1. *Measurement coarse-graining is unavoidable.* Finite distinguishability implies no arbitrarily sharp POVMs exist in bounded regions. At some scale, measurement precision must saturate. This is consistent with quantum uncertainty relations and suggests operational limits on detector resolution.
2. *Holographic information bounds.* If finite distinguishability connects to the Bekenstein bound (Appendix D.2), then information storage in any region is bounded by surface area, not volume. This is testable in principle through black hole thermodynamics and holographic systems.
3. *Discrete vs. continuous structure.* The framework is compatible with both continuous and discrete time (see Remark in §5.2.1). Observations distinguishing these cases (e.g., Planck-scale discreteness) would refine but not refute the framework.

What would genuinely falsify the framework:

- Demonstrated infinite operational distinguishability within bounded systems
- Systematic, reproducible deviations from Born-rule probabilities
- Operational contextuality (different statistics for operationally identical procedures)
- Faster-than-light signaling via entanglement

None of these has been observed. The framework's structural predictions match a century of experimental quantum mechanics.

11.11 What Would It Mean for Admissibility to Fail?

The admissibility framework introduced in this work is intentionally minimal, yet it imposes strong structural constraints. It is therefore important to clarify what it would mean, physically and conceptually, for admissibility to fail.

In Plain Language: We've argued that admissibility is required for facts to exist. But what if we're wrong? This subsection examines each way admissibility could fail and shows that each failure mode would undermine not just quantum mechanics, but the very possibility of doing physics at all.

Admissibility would fail if one or more of the following conditions were violated:

Infinite operational distinguishability: If arbitrarily fine distinctions were physically accessible within bounded systems, then any apparent irreversible process could be refined into a reversible one. In such a universe, no outcome would ever be final; all records would remain recoverable in principle. Facts would be provisional rather than stable. The very notion of an experimental result would lose operational meaning.

Absence of irreversible commitment: If no admissible process could irreversibly commit distinctions—if every physical transformation were invertible within the admissible domain—then records, memories, and measurements would not exist as physical facts. Physics would reduce to a purely reversible formal system with no mechanism for outcome fixation.

Unbounded physical resources: If physical processes could exploit infinite energy, infinite memory, or infinite precision, then the operational constraints defining admissibility would be void. However, such assumptions are incompatible with the finite localization, finite control, and finite duration of real physical processes, and they undermine the empirical basis of physics itself.

Context-dependent probability assignments: If outcome probabilities depended on physically inaccessible implementation details—beyond what admissible procedures can resolve—then probability would lose operational meaning. Identical experimental situations could yield incompatible statistical descriptions, contradicting the reproducibility of empirical science.

A failure of admissibility in any of these senses would not merely modify the conclusions of this paper; it would call into question the possibility of physics as an empirical discipline. The admissibility framework does not assume the existence of facts as a metaphysical postulate. Rather, it treats the existence of facts as a minimal operational requirement for doing physics at all.

In Plain Language: Admissibility isn't one assumption among many that might turn out to be false. It's the condition for there being anything to be true or false *about*. A universe without admissibility wouldn't be a different kind of physics—it would be a universe where physics, as an empirical enterprise, couldn't exist.

From this perspective, admissibility is not an optional interpretive stance. It is a consistency condition on any framework that purports to describe a world in which experiments can be performed, outcomes recorded, and results compared.

---

12. A Skeptic's FAQ

This section addresses the most common—and most serious—objections to the admissibility framework. These are not strawman questions; they are the exact concerns a careful skeptic, physicist, or philosopher should raise.

"Aren't you just assuming quantum mechanics in disguise?"

Short answer: No. The paper does not assume Hilbert space, unitarity, or probabilities. It derives them.

Longer answer: The framework begins *below* quantum mechanics, at the level of what any empirical physics must support: facts, records, and irreversible outcomes. From this starting point, we derive constraints (finite distinguishability and irreversible commitment) that any fact-producing theory must satisfy.

Only after these constraints are established do we introduce operational postulates governing reversible dynamics (convexity, continuity, compositionality). These postulates are not quantum axioms; they are requirements for reversible evolution to coexist consistently with irreversible fact-production.

Hilbert space appears only at the end, as the unique mathematical representation compatible with these constraints. It is a result, not a premise.

"Why should physics care about operational limits? What about ontology?"

Short answer: Physics is an empirical discipline. If a distinction is operationally inaccessible, it cannot ground a physical fact.

Longer answer: The admissibility framework is explicitly operational. It makes no claims about what "exists in itself" beyond what can be accessed, distinguished, and recorded using finite physical procedures.

If two states differ ontologically but agree on all admissible measurements, they are empirically identical. Treating such differences as physically real is a metaphysical choice, not a physical one.

The paper's results constrain operational physics—the domain where experiments, records, and facts live. Ontological excess that cannot be operationally accessed is outside the scope of physics as an empirical science.

"Isn't finite distinguishability just a limitation of our instruments?"

Short answer: No. Finite distinguishability is *required* for irreversible facts to exist at all.

Longer answer: The paper proves a no-go theorem: if a system permits infinite operational distinguishability, then no process can irreversibly commit distinctions. Any apparent loss of information can always be recovered by accessing finer degrees of freedom.

But facts—measurements, records, memories—require irreversible loss of alternatives. Therefore, if facts exist, operational distinguishability must be finite.

This is not about imperfect instruments. It is about the logical requirements for fact-production in any physically realizable universe.

"Doesn't classical chaos already give irreversibility without quantum mechanics?"

Short answer: Only by assuming finite distinguishability.

Longer answer: Classical chaos amplifies microscopic differences into macroscopic outcomes—but only once those outcomes are coarse-grained into finite records. The moment a chaotic system produces a stable fact (a bit, a symbol, a pointer position), continuous degrees of freedom are compressed into discrete outcomes.

That compression is exactly what the admissibility framework analyzes. Classical statistical mechanics achieves irreversibility by assuming operational inaccessibility of microscopic details. The framework explains why that assumption is necessary and why it cannot be removed at a fundamental level.

"Doesn't Many-Worlds avoid irreversible commitment entirely?"

Short answer: No—only at the ontological level, not the operational one.

Longer answer: Everettian interpretations deny fundamental collapse, but observers within a branch still experience definite outcomes, form records, and accumulate entropy. From the perspective of any observer doing physics, facts are irreversible relative to accessible degrees of freedom.

The admissibility framework operates at this operational level. It does not rule out Many-Worlds, but it shows that even Everettian physics must obey finite distinguishability, CPTP dynamics, and Born-rule statistics as experienced.

Interpretations differ on ontology; admissibility constrains operational structure.

"Aren't you assuming non-contextuality to get the Born rule?"

Short answer: No. Operational non-contextuality is *forced* by finite distinguishability.

Longer answer: If the probability of an outcome depended on which other outcomes were possible but not realized, that dependence would require some physical degree of freedom encoding the measurement context.

But if that degree of freedom were operationally accessible, it would appear in the outcome record—making the contexts distinguishable. If it were inaccessible, it could not influence observable probabilities.

Therefore, outcome probabilities must depend only on the state and the effect that actually occurred. This is operational non-contextuality, not an ontological assumption.

Given operational non-contextuality, the Born rule follows uniquely via Gleason's theorem.

"Why do you need tomographic locality? Isn't that an extra assumption?"

Short answer: Tomographic locality is required for all physically accessible information to be measurable.

Longer answer: If a composite system had physically relevant degrees of freedom that were not reflected in any local or correlational measurements, those degrees of freedom would be operationally inaccessible.

The admissibility framework treats operationally inaccessible structure as physically irrelevant. Physics describes what can, in principle, be measured and recorded. Tomographic locality formalizes this requirement for composite systems.

Relaxing tomographic locality admits theories with hidden global structure—but such structure cannot ground observable facts.

"Why continuous reversible dynamics? Why not discrete jumps?"

Short answer: Discontinuous reversible dynamics are empirically indistinguishable from stochastic commitment under finite control.

Longer answer: If infinitesimal changes in experimental control produced discontinuous changes in system evolution, those discontinuities would function as facts: stable, uncontrollable distinctions.

To prevent reversible evolution from itself generating commitments, reversible dynamics must vary continuously with control parameters. Continuity is therefore required to separate reversible evolution from irreversible fact-production.

"Why complex Hilbert space specifically? Why not real or quaternionic?"

Short answer: Because only complex quantum mechanics satisfies all admissibility-motivated requirements simultaneously.

Longer answer: Real and quaternionic quantum theories fail at least one of the following:

- Continuous transitive symmetry on pure states
- Tomographic locality for composites
- Full interference structure
- Consistent composition of subsystems

These failures introduce either operationally inaccessible degrees of freedom or discontinuities in reversible dynamics—both incompatible with admissibility.

Complex Hilbert space is not chosen; it is what remains after all admissibility-violating alternatives are excluded.

"What does this framework NOT explain?"

Short answer: It explains structure, not parameters.

Longer answer: The admissibility framework explains why quantum mechanics has:

- Unitarity
- Hamiltonians
- The Born rule
- Entanglement
- Finite outcomes
- Irreversible measurement

It does *not* explain:

- Specific interaction Hamiltonians
- Coupling constants
- Particle masses
- Number of generations
- Gravity

Those require additional physical principles. The claim here is not "quantum mechanics explains everything," but rather: quantum mechanics is the only *framework* compatible with the existence of facts.

"What about concrete alternative theories—Spekkens' toy model, PR boxes, etc.?"

Short answer: They all fail admissibility in specific ways.

Longer answer: Several explicit non-quantum theories have been constructed:

Spekkens' toy model [20] is an epistemic model where "quantum" states represent knowledge about underlying classical states. It reproduces many quantum phenomena but is fundamentally epistemic—it doesn't support genuine irreversible commitment because the underlying classical states are always well-defined. In Spekkens' model, "measurement" updates our knowledge about a pre-existing ontic state; the epistemic restriction limits what we can *know*, not what *exists*. This is knowledge-update, not physical commitment, so the model illustrates epistemic restrictions rather than fact-producing physics.

PR boxes are hypothetical devices producing maximally nonlocal correlations (violating Bell inequalities more strongly than quantum mechanics). They satisfy no-signaling but violate information-theoretic principles like "information causality" (Pawłowski et al., 2009) that follow from the Born-rule structure derived in Theorem 7.1. From an admissibility perspective, PR boxes would require correlations that exceed what finite-dimensional Hilbert spaces can support—they're incompatible with the Jordan-algebraic structure forced by R1–R5.

Classical stochastic theories with genuine randomness face the problem that randomness without finite distinguishability doesn't produce stable facts—the outcomes keep refining indefinitely. If you impose finite distinguishability, you're back to quantum structure.

These examples illustrate that admissibility constraints have teeth: proposed alternatives fail in identifiable ways.

Bottom Line for Skeptics

You do not have to agree with the interpretation. You do not have to like the conclusion.

But to reject the framework, you must identify one of the following:

1. A way to have irreversible facts with infinite operational distinguishability, or
2. A way for inaccessible degrees of freedom to influence observable statistics, or
3. A reversible dynamics that produces no commitments despite discontinuities, or
4. An alternative mathematical structure that satisfies all admissibility constraints simultaneously.

Absent that, quantum mechanics is not a mysterious option among many—it is the unique survivor of the requirements for a factual universe.

---

13. Conclusion

We have shown that quantum mechanics is not an arbitrary theoretical invention, nor a contingent framework selected solely by empirical success. Rather, it emerges as the unique structural framework compatible with physical admissibility in a universe capable of producing irreversible facts.

The results fall into two tiers. **Tier I** (forced by admissibility): any fact-producing physics must exhibit finite distinguishability, irreversible commitment, and distinguishability-preserving reversible dynamics. These constraints exclude infinite-precision classical mechanics and carve out a narrow class of admissible theories. **Tier II** (selection within that class): operational closure principles—tomographic locality and maximal reversibility—uniquely select complex Hilbert-space quantum mechanics from among the admissible options.

Starting from the minimal requirement that physical records exist, we identified the unavoidable constraints and traced their consequences. Reversible evolution must be unitary and Hamiltonian-generated. Measurement arises as minimal irreversible extension. Probability is fixed uniquely by the Born rule. Entanglement and Bell correlations follow from global commitment structure. Relativistic consistency forces spinorial dynamics.

These results relocate quantum foundations from postulates to necessity. The familiar mathematical structures are no longer assumed but explained. Quantum mechanics appears not as a mysterious departure from classical reasoning, but as the inevitable architecture of any fact-producing physical world.

Final Thought for the General Reader: For a century, we've asked "why is quantum mechanics true?" as if some other physics might have been possible. This paper suggests a different answer: quantum mechanics is the unique structural framework for fact-production. The question isn't why quantum mechanics—it's how facts are possible, and quantum mechanics is the answer.

In this sense, quantum mechanics does not require interpretation to be justified. What requires justification is the assumption that stable physical records could exist in a universe governed by anything else.

What This Changes

This section is intended for general readers and summarizes the broader implications of the admissibility framework.

The measurement problem looks different now. For decades, physicists have asked: "Why does measurement cause collapse?" The admissibility framework suggests this question conflates two issues. The first—why measurement has the *structure* it does (irreversible commitment, finite outcomes, Born-rule probabilities)—is answered here: that structure is required for facts to exist. The second—why a *particular* outcome is selected—remains open, a question for interpretation rather than architecture. Collapse isn't a problem to be solved; it's commitment doing exactly what commitment does. What remains unexplained is why *this* outcome rather than that one, and that question may have no further physical answer.

Interpretation debates are downstream. Copenhagen, Many-Worlds, Bohmian mechanics, QBism—these interpretations disagree about what the quantum state *is*, but they all use the same mathematical structure. The admissibility framework explains why that structure is non-negotiable: it's forced by the requirement that facts exist. Interpretation is about the meaning of a structure that was never optional in the first place. This doesn't resolve interpretation debates, but it reframes them: they're about semantics, not architecture.

Classical physics is the special case, not the default. We're taught classical mechanics first, then told quantum mechanics is "weird" and needs explaining. But the admissibility analysis reverses this. Classical mechanics assumes infinite distinguishability and reversible determinism—both incompatible with irreversible facts. Classicality is what emerges when quantum systems decohere and commit at scales where the underlying quantum structure is operationally inaccessible. Classical physics is the *derived* approximation; quantum mechanics is the *necessary* foundation.

What's next: from form to content. This paper explains why quantum mechanics has the *structure* it does—unitarity, Hamiltonians, Born rule, entanglement. It doesn't explain the *content*: why the fine-structure constant is 1/137, why there are three fermion generations, why gravity is weak. That's where the broader VERSF program comes in. The Void Energy-Regulated Space Framework proposes that even these "free parameters" emerge from deeper admissibility-like constraints—specifically, from information-geometric requirements on how entropy, distinguishability, and commitment structure can be embedded in spacetime. If that program succeeds, the arbitrary constants of physics would be as necessary as the quantum formalism itself.

The present paper is the first step: establishing that the *framework* is forced. The next step is showing that the *parameters* are too.

References

[1] L. Hardy, "Quantum Theory From Five Reasonable Axioms," arXiv:quant-ph/0101012 (2001).

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

[3] L. Masanes and M. P. Müller, "A derivation of quantum theory from physical requirements," New J. Phys. **13**, 063001 (2011).

[4] B. Dakić and Č. Brukner, "Quantum Theory and Beyond: Is Entanglement Special?" in *Deep Beauty*, ed. H. Halvorson (Cambridge University Press, 2011).

[5] J. Barrett, "Information processing in generalized probabilistic theories," Phys. Rev. A **75**, 032304 (2007).

[6] W. H. Zurek, "Decoherence, einselection, and the quantum origins of the classical," Rev. Mod. Phys. **75**, 715 (2003).

[7] C. W. Helstrom, *Quantum Detection and Estimation Theory* (Academic Press, 1976).

[8] H. Barnum, M. P. Müller, and C. Ududec, "Higher-order interference and single-system postulates characterizing quantum theory," New J. Phys. **16**, 123029 (2014).

[9] M. P. Müller and L. Masanes, "Information-theoretic postulates for quantum theory," in *Quantum Theory: Informational Foundations and Foils*, eds. G. Chiribella and R. W. Spekkens (Springer, 2016).

[10] M. Koecher, "Die Geodätischen von Positivitätsbereichen," Math. Ann. **135**, 192 (1958); E. B. Vinberg, "Homogeneous cones," Soviet Math. Dokl. **1**, 787 (1960).

[11] P. Jordan, J. von Neumann, and E. Wigner, "On an Algebraic Generalization of the Quantum Mechanical Formalism," Ann. Math. **35**, 29 (1934).

[12] E. P. Wigner, "On the Quantum Correction For Thermodynamic Equilibrium," Phys. Rev. **40**, 749 (1932); E. P. Wigner, *Group Theory and its Application to Quantum Mechanics* (Academic Press, 1959).

[13] V. Bargmann, "On Unitary Ray Representations of Continuous Groups," Ann. Math. **59**, 1 (1954).

[14] M. H. Stone, "On One-Parameter Unitary Groups in Hilbert Space," Ann. Math. **33**, 643 (1932).

[15] M.-D. Choi, "Completely positive linear maps on complex matrices," Lin. Alg. Appl. **10**, 285 (1975).

[16] K. Kraus, *States, Effects, and Operations: Fundamental Notions of Quantum Theory* (Springer, 1983).

[17] A. M. Gleason, "Measures on the Closed Subspaces of a Hilbert Space," J. Math. Mech. **6**, 885 (1957).

[18] P. Busch, "Quantum States and Generalized Observables: A Simple Proof of Gleason's Theorem," Phys. Rev. Lett. **91**, 120403 (2003).

[19] C. M. Caves, C. A. Fuchs, K. K. Manne, and J. M. Renes, "Gleason-Type Derivations of the Quantum Probability Rule for Generalized Measurements," Found. Phys. **34**, 193 (2004).

[20] R. W. Spekkens, "Contextuality for preparations, transformations, and unsharp measurements," Phys. Rev. A **71**, 052108 (2005).

[21] J. S. Bell, "On the Einstein Podolsky Rosen Paradox," Physics **1**, 195 (1964).

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

[23] R. Bousso, "The holographic principle," Rev. Mod. Phys. **74**, 825 (2002).

[24] S. L. Adler, *Quaternionic Quantum Mechanics and Quantum Fields* (Oxford University Press, 1995); A. Peres, "Proposed Test for Complex versus Quaternion Quantum Theory," Phys. Rev. Lett. **42**, 683 (1979).

Appendix A: Assumptions and Regularity Conditions

This appendix lists explicitly the assumptions and regularity conditions employed in the main text. The purpose is not to weaken the results, but to make transparent which elements are taken as minimal physical requirements and which serve as technical regularity conditions enabling mathematical representation.

A.1 Minimal Physical Assumptions

The following assumptions are treated as methodological preconditions for empirical physics:

(A1) Operational Recordability Empirical physics presupposes the existence of stable, communicable records ("operational facts") produced by physical processes. This is not asserted as a claim about fundamental ontology; it is a methodological requirement for any framework that purports to be testable, comparable across observers, and usable for inference from experiments.

Scope and status. We do not claim that fact-production is ontologically fundamental. Record formation may be emergent rather than primitive. However, even if one holds that records emerge from a more fundamental reversible dynamics, the admissibility constraints apply to the operational layer at which any empirical theory is validated. The results therefore constrain any viable fundamental theory *indirectly*, via the requirement that it recover a record-forming operational regime.

In other words: wherever stable records appear—whether fundamentally or emergently—the constraints of Sections 3–7 apply. A fundamental theory that cannot recover an operational layer with finite distinguishability and irreversible commitment cannot be empirically tested, and therefore falls outside the scope of physics as an experimental science.

(A2) Finite operational resources Physical processes are implemented using finite energy, finite time, finite spatial extent, finite memory, and finite resolution. This excludes procedures requiring infinite precision or unbounded refinement.

(A3) Physical realizability of measurements If two states are operationally distinguishable, there exists an admissible physical procedure that distinguishes them with finite resources.

These assumptions are not specific to quantum mechanics and are satisfied by any empirically meaningful physical theory.

A.2 Admissibility Constraints

From the minimal assumptions above, the following admissibility constraints are defined:

(C1) Finite distinguishability In any bounded physical system, there exists a finite bound on the number of mutually distinguishable states accessible under admissible procedures.

(C2) Irreversible commitment There exist admissible processes that map multiple prior possibilities to a single outcome in a way that cannot be undone by any admissible recovery procedure.

These constraints are the core structural inputs to the results of this paper.

A.3 Operational Postulates for Reversible Dynamics

The following postulates govern reversible dynamics between commitments. They are stated in full in Section 5.2.1 and summarized here with their admissibility status. Crucially, R1, R4, and R5 are *derived* from minimal experimental capability assumptions; see Appendix E for full proofs.

(R1) Convexity (Operational Mixing) Classical randomness over preparations yields convex mixtures. *Status:* Derived from classical control capability (Assumption CC, Lemma E.1).

(R2) Continuous Reversible Dynamics Reversible evolution forms a continuous, connected group. *Status:* Forced by admissibility; discontinuous dynamics would generate commitments.

(R3) Distinguishability-Preserving Reversibility Reversible transformations preserve operational distinguishability. *Status:* Forced by admissibility; erasing distinctions is commitment; amplifying them creates facts.

(R4) Operational Closure of Composition Operationally indistinguishable composite states are identified; joint states are determined by local measurements and correlations. *Status:* Derived from universal composite controllability (Assumption UC, Theorem E.5).

(R5) Maximal Reversibility (Purification) Mixed states arise as marginals of pure states; purifications are equivalent up to local unitaries. *Status:* Derived from reversible embeddability (Assumption RE, Theorem E.3).

Postulate classification: R2 and R3 are directly forced by admissibility. R1, R4, and R5 are derived from assumptions about experimental capabilities (CC, UC, RE) that any laboratory doing physics must have. These are not quantum-specific axioms but minimal operational completeness requirements.

Under R1–R5, complex Hilbert space quantum mechanics is uniquely selected (Lemma 5.1).

A.4 Status of the Results

The central conclusions of the paper—finite distinguishability, necessity of unitary dynamics, existence of the Hamiltonian, measurement as irreversible commitment, uniqueness of the Born rule, and admissibility-fixed-point characterization—depend on assumptions (A1)–(A3), constraints (C1)–(C2), and derived postulates (R1)–(R5).

The logical structure is:

- (A1)–(A3) are preconditions for empirical physics
- (C1)–(C2) follow from (A1)–(A3) via Theorem 3.2
- (CC), (UC), (RE) are minimal experimental capability assumptions
- (R1), (R4), (R5) are derived from (CC), (UC), (RE) via Appendix E
- (R2), (R3) are forced directly by (C1)–(C2)
- Hilbert space structure follows from (R1)–(R5) via Lemma 5.1

The experimental capability assumptions (CC, UC, RE) express capabilities any laboratory doing physics must have: classical control, access to extended systems, and universal operations on composites. They are not quantum-specific but apply to any operational theory claiming to describe experimentally accessible degrees of freedom.

In Plain Language: This appendix separates what we *must* assume (that facts exist, resources are finite, measurements are possible, labs have basic capabilities) from what we *derive* (convexity, tomographic locality, purification). The postulates R1, R4, R5 are no longer axioms about nature—they're theorems about what follows from having a functional laboratory.

Appendix B: Extended Proofs and Technical Details

This appendix provides complete proofs for results given as sketches in the main text, addresses technical subtleties, and clarifies the ontological versus operational status of key claims.

B.1 Ontological vs. Operational Clarification

Throughout this paper, we maintain a strict operational stance: claims are about what can be *physically accessed, distinguished, and recorded* using finite procedures, not about what "really exists" at some deeper ontological level. This distinction is crucial for interpreting the main results correctly.

The operational interpretation of finite distinguishability: When we assert that a system has finite distinguishability N, we do not claim that "only N states exist." We claim that no admissible procedure—no physical measurement implementable with finite resources—can reliably distinguish more than N possibilities. The underlying ontology may be richer; what matters is operational accessibility.

The operational interpretation of irreversible commitment: When we say a process is irreversibly committing, we do not claim information is "destroyed" in some absolute sense. We claim it becomes inaccessible to all admissible recovery procedures. In an Everettian interpretation, information may persist in other branches; operationally, it is gone from the accessible record.

Why this matters for Theorem 3.2: The theorem shows that infinite *operational* distinguishability is incompatible with irreversible commitment. A skeptic might object: "What if infinitely many states exist ontologically, but only finitely many are operationally accessible?" This objection actually supports our framework. If operational access is finite while ontology is infinite, then the operational physics—the physics that produces facts—satisfies finite distinguishability. The theorem applies to operational physics, which is what empirical science studies. Ontological claims beyond operational access are metaphysics, not physics.

B.2 Full Proof of Theorem 3.2 (No-Go for Infinite Distinguishability)

We provide a fully rigorous version of the proof with explicit base case and inductive structure.

Setup. Let $\mathcal{S}$ be a state space with infinite operational distinguishability in the sense of Definition 3.1: for every $\varepsilon > 0$ and every pair of distinct states $\rho_1 \neq \rho_2$, there exists an admissible discrimination procedure $D\varepsilon$ achieving error probability less than ε.

Claim. No admissible process $\Phi: \mathcal{S} \to \mathcal{S}$ can exhibit irreversible commitment.

Proof.

Base case: Consider any admissible process Φ and suppose $\Phi(\rho_1) = \Phi(\rho_2) = \rho^*$ for distinct states $\rho_1 \neq \rho_2$. We show that the distinction between ρ_1 and ρ_2 can be recovered after Φ.

By Definition 3.1, for $\varepsilon = 1/4$, there exists an admissible procedure $D_{1/4}$ that distinguishes ρ_1 from ρ_2 with error < 1/4. This procedure produces a classical output $c \in \{1, 2\}$ satisfying:

- $P(c = 1 \mid \rho_1) > 3/4$
- $P(c = 2 \mid \rho_2) > 3/4$

Inductive construction: Consider the composite procedure: first apply $D_{1/4}$ to produce record c, then apply Φ to the quantum system. The joint state of (record, quantum system) is:

- Starting from ρ_1: $(c \approx 1, \rho^*)$ with probability > 3/4
- Starting from ρ_2: $(c \approx 2, \rho^*)$ with probability > 3/4

The record c survives the action of Φ on the quantum system (records are stable by definition of commitment). Therefore, after Φ acts, reading c reveals whether the original state was ρ_1 or ρ_2 with error < 1/4.

Amplification: By repeating with $\varepsilon = 1/4^n$ for n = 1, 2, 3, ..., we obtain a sequence of discrimination procedures with exponentially decreasing error. For any desired confidence level $1 - \delta$, there exists an admissible procedure that recovers the ρ_1-ρ_2 distinction with error $< \delta$.

Conclusion: Since the distinction can be recovered with arbitrary accuracy using admissible procedures, Φ does not irreversibly commit the ρ_1-ρ_2 distinction. Since ρ_1, ρ_2 were arbitrary distinct states mapped to the same output, Φ exhibits no irreversible commitment. Since Φ was arbitrary, no admissible process exhibits irreversible commitment under infinite distinguishability. ■

B.3 Self-Duality as Definitional Closure

A potential objection: "Measurement completeness is itself a postulate, so deriving self-duality from it just moves the assumption." This objection misunderstands the logical status of the derivation.

Self-duality is not derived from quantum mechanics; it is a *methodological closure condition* on what counts as a state or effect in an operational theory:

- **States** are equivalence classes of preparation procedures, identified when they produce identical measurement statistics
- **Effects** are equivalence classes of measurement outcomes, identified when they respond identically to all states

If a mathematical effect E is compatible with all states (gives valid probabilities) but has no physical implementation, then E defines no operational distinction. Such effects should be quotiented out of the theory. Similarly, if two states agree on all implementable measurements, they are operationally identical and should be identified.

Formal statement. Let $\mathcal{S}$ be the operational state space and $\mathcal{S}^*$ the cone of all linear functionals that are non-negative on states and bounded by 1. *Measurement completeness* asserts $\mathcal{S}^* = \{\text{implementable effects}\}$. *State completeness* asserts that states differing only on non-implementable effects are identified.

Derivation of self-duality. Under measurement completeness and state completeness:

1. Every state ρ defines a functional on effects via $E \mapsto \mathrm{Tr}(E\rho)$, so $\mathcal{S}$ embeds in $\mathcal{S}^*$
2. Measurement completeness implies every element of $\mathcal{S}^*$ is realized by some effect
3. State completeness implies the embedding is injective (distinct states differ on some effect)
4. Finite distinguishability implies finite dimensionality, so the embedding is surjective

Therefore $\mathcal{S} \cong \mathcal{S}^*$, which is self-duality.

Why this isn't circular. The argument doesn't assume quantum mechanics. It defines what "state" and "effect" mean in any operational theory: they are equivalence classes under operational indistinguishability. Self-duality then follows as a consistency condition. A theory violating self-duality would have either:

- Effects with no physical implementation (metaphysical, not physical)
- States distinguishable only by non-implementable effects (operationally identical, should be identified)

Neither possibility describes a coherent operational physics.

B.4 Full Proof of Lemma 5.1 (Hilbert Space Emergence)

Given: A convex state space $\mathcal{S}$ satisfying:

- (i) Finite distinguishability (dimension N)
- (ii) Continuous transitive symmetry group G of distinguishability-preserving transformations
- (iii) Self-duality

To prove: $\mathcal{S}$ is isomorphic to density matrices on $\mathbb{C}^N$.

Step 1: Jordan algebra structure. By the Koecher-Vinberg theorem, a finite-dimensional convex cone is self-dual and homogeneous (admits a transitive automorphism group) if and only if it is the cone of squares $C = \{x^2 : x \in A\}$ in a formally real Jordan algebra A. Here, a Jordan algebra satisfies $x \circ y = y \circ x$ and $x \circ (y \circ x^2) = (x \circ y) \circ x^2$, and "formally real" means $x^2 + y^2 = 0$ implies $x = y = 0$.

Step 2: Classification. The Jordan-von Neumann-Wigner theorem classifies simple finite-dimensional formally real Jordan algebras:

- Type $I_n(\mathbb{R})$: n×n symmetric real matrices, Jordan product $A \circ B = ½(AB + BA)$
- Type $I_n(\mathbb{C})$: n×n Hermitian complex matrices
- Type $I_n(\mathbb{H})$: n×n Hermitian quaternionic matrices
- Spin factors: $\mathbb{R} \oplus \mathbb{R}^d$ with product $(\alpha, v) \circ (\beta, w) = (\alpha\beta + v \cdot w, \alpha w + \beta v)$
- Type $I_3(\mathbb{O})$: 3×3 Hermitian octonionic matrices (the Albert algebra, 27-dimensional)

Step 3: Exclusion of Albert algebra. The Albert algebra lacks an associative envelope: it cannot be embedded in any associative algebra of matrices. This has physical consequences:

- No consistent tensor product: given two Albert-algebra systems A and B, there is no natural Albert-algebra structure on $A \otimes B$ that satisfies local tomography (the joint state is determined by correlations of local measurements)
- No consistent dynamics: the symmetry group F_4 of the Albert algebra is exceptional and does not embed into $GL(n, \mathbb{C})$ in a way compatible with continuous composition

Compositional locality (R4) requires tensor products; continuous composability (R2) requires embeddability into continuous groups. The Albert algebra satisfies neither.

Step 4: Exclusion of quaternionic quantum mechanics. Quaternionic quantum mechanics (QQM) on $\mathbb{H}^n$ has state space = density matrices over $\mathbb{H}^n$. For $n \geq 2$, QQM differs from complex QM (CQM) observationally:

Local tomography failure: In QQM, the joint state of two systems is not determined by the outcomes of all local measurements and their correlations. There exist distinct joint states $\rho_{12} \neq \sigma_{12}$ such that $Tr[(A \otimes B)\rho_{12}] = Tr[(A \otimes B)\sigma_{12}]$ for all local observables A, B. This means operational access to the full state space requires nonlocal measurements, violating the spirit of compositional locality.

Bit asymmetry: In QQM, not all pairs of orthogonal pure states are equivalent under the symmetry group Sp(n). The "equator" of the state space has different structure than the "poles." Continuous transitivity requires that any pure state can be continuously connected to any other via reversible operations; in QQM with $n \geq 3$, certain transformations require discontinuous jumps.

The decisive argument is local tomography [3]: CQM is the unique theory in the Jordan classification satisfying local tomography for all composite systems.

Step 5: Exclusion of real quantum mechanics. Real quantum mechanics (RQM) on $\mathbb{R}^n$ uses real symmetric matrices. Failures:

Bit asymmetry: For a two-level system, RQM gives a 2-dimensional Bloch disk rather than the 3-dimensional Bloch ball of complex QM. Pure states lie on a circle S^1 rather than a sphere S^2.

Antipodal pure states cannot be connected by continuous reversible transformations (the group O(2) is disconnected). This violates transitivity.

Interference asymmetry: In CQM, interference between any two paths is governed by a phase e^{iθ}; relative phases can be continuously varied. In RQM, the "phase" is ±1, and interference is either constructive or destructive with no intermediate values. This discreteness conflicts with continuous composability for certain process combinations.

Local tomography: RQM also fails local tomography for certain composite systems.

Step 6: Conclusion. The unique Jordan algebra satisfying self-duality, continuous transitive symmetry, compositional tensor products, and local tomography is Type $I_n(\mathbb{C})$—Hermitian matrices over $\mathbb{C}$. The state space is density matrices on $\mathbb{C}^n$, i.e., quantum mechanics. ■

B.5 Contextuality and Spekkens' Framework

The argument that admissibility forces non-contextuality deserves deeper engagement with Spekkens' work on operational and ontological contextuality [20].

Spekkens' distinction. Spekkens distinguishes:

- *Preparation contextuality:* The same quantum state ρ may arise from different preparation procedures; if the ontic state depends on which procedure was used, preparations are ontologically contextual
- *Measurement contextuality:* The same POVM element E_i may appear in different POVMs; if the ontic response depends on which POVM contains E_i, measurements are ontologically contextual
- *Transformation contextuality:* Similar for channels

Our claim (refined). Admissibility forces *operational* non-contextuality: outcome statistics cannot depend on operationally inaccessible features. This is weaker than claiming ontological non-contextuality.

Detailed argument. Suppose measurement outcome probabilities were operationally contextual: $p(E_i|\rho, POVM_1) \neq p(E_i|\rho, POVM_2)$ where E_i is the same effect in both POVMs and ρ is the same state.

This difference in probability must be detectable in outcome statistics. Let f_1 = observed frequency of outcome i in $POVM_1$, f_2 = observed frequency in $POVM_2$. If $p(...|POVM_1) \neq p(...|POVM_2)$, then f_1 and f_2 will systematically differ.

But the only record produced is "outcome i occurred." The record does not encode which POVM was implemented (that would require additional degrees of freedom in the record, making the POVMs operationally distinguishable). Under finite distinguishability, the observer cannot access which POVM context applies. Therefore, any statistical difference $f_1 \neq f_2$ would be a systematic reproducible difference with no physically accessible cause—violating the principle that operational distinctions require physical distinctions.

Compatibility with ontological contextuality. Our argument does not rule out ontological contextuality in Spekkens' sense. Hidden variables might respond differently to different contexts. But if this contextuality is operationally invisible (averages out, or is encoded in inaccessible degrees of freedom), it is compatible with admissibility. What's ruled out is operationally manifest contextuality—different outcome statistics for the same operational procedure.

Could contextual theories satisfy admissibility? Yes, if their contextuality is ontological rather than operational. Bohmian mechanics, for instance, is preparation-contextual (the particle position depends on how ρ was prepared) but operationally reproduces quantum statistics. Such theories satisfy admissibility at the operational level while differing ontologically.

B.6 Relation to Decoherence and Einselection

A natural question: Is "irreversible commitment" just decoherence under a different name?

Decoherence (Zurek's einselection) describes how quantum coherence is suppressed when a system interacts with an environment. The reduced density matrix of the system rapidly becomes diagonal in a preferred "pointer basis" determined by the system-environment interaction.

Irreversible commitment as defined here is the production of a stable record that cannot be undone by admissible procedures.

Key differences:

1. *Decoherence is continuous; commitment is discrete.* Decoherence describes a continuous loss of coherence over time. Commitment produces a discrete fact—a definite outcome from a finite set. Decoherence alone doesn't explain why outcomes are discrete (finite distinguishability does) or why specific outcomes occur (the selection problem).
2. *Decoherence is in principle reversible; commitment is not.* Decoherence transfers coherence to the environment but doesn't destroy it. In principle, by controlling all environmental degrees of freedom, coherence could be recovered. Commitment, by contrast, is irreversible *within the admissible domain*. The difference is operational: admissibility bounds what procedures can access.
3. *Decoherence doesn't select outcomes; commitment does.* Decoherence explains why interference is suppressed and why the density matrix becomes diagonal. It doesn't explain why a particular diagonal element is realized as fact. Commitment is the realization of fact; decoherence prepares the stage.

Connection: Decoherence is a *mechanism* that produces the conditions for commitment. In many physical situations, the system-environment interaction rapidly decoheres the system, and then the environment serves as the "record" that commits the outcome. But commitment is the conceptual category; decoherence is one (important, ubiquitous) physical mechanism that enables it.

Pointer basis and admissibility: Zurek's einselection derives the pointer basis from the criterion that pointer states are stable under environmental monitoring. This resonates with admissibility: states that survive as facts must be robust against the environmental "measurements" that constitute ongoing physical interaction. The admissibility framework provides a structural explanation for why such preferred bases exist (they're the eigenbases of commitment operations), while decoherence explains how the selection is implemented dynamically.

B.7 Relation to the PBR Theorem

The Pusey-Barrett-Rudolph (PBR) theorem [25] shows that if quantum states represent states of reality (ψ-ontic) rather than merely states of knowledge (ψ-epistemic), then certain independence assumptions hold. The theorem rules out a class of ψ-epistemic hidden variable models.

Connection to admissibility: The admissibility framework is operationally neutral about whether ψ is ontic or epistemic—it constrains what's operationally accessible, not what "really exists." However, the structural results have implications:

1. *Operational uniqueness implies representational uniqueness.* If quantum states are uniquely determined by their operational statistics (as in our framework), and if operational statistics are what physical theories should capture, then there's no room for multiple ontic states underlying the same quantum state. This favors ψ-ontic interpretations, though doesn't strictly require them.
2. *Finite distinguishability constrains hidden variables.* Any hidden variable theory must respect finite distinguishability at the operational level. This means hidden variables cannot be operationally accessed beyond the quantum bound N. ψ-epistemic models that require operationally accessible hidden structure are excluded.

Our position: The admissibility framework is consistent with both ψ-ontic and ψ-epistemic interpretations, provided the latter don't posit operationally accessible structure beyond quantum states. The PBR theorem, combined with admissibility, significantly constrains viable ψ-epistemic models.

B.8 Connection to Thermal Time and the Connes-Rovelli Hypothesis

Proposition 4.4 derives temporal ordering from commitment structure: time's direction is the direction of increasing commitment and entropy. This resonates with the Connes-Rovelli thermal time hypothesis [26].

Thermal time hypothesis: In generally covariant theories without preferred time, Connes and Rovelli propose that time emerges from the state of the system via the modular flow of the observable algebra. Given a state ω, the Tomita-Takesaki theorem provides a one-parameter group of automorphisms—the "thermal time" flow—relative to which ω appears thermal.

Comparison:

Aspect	Admissibility (this paper)	Thermal Time (Connes-Rovelli)
Time from	Commitment/entropy increase	Modular flow of state
Fundamental object	Facts/records	Algebraic state
Arrow of time	Commitment ordering	Thermal equilibrium
Mechanism	Information erasure	KMS condition

Potential synthesis: Both approaches derive temporal structure from state-dependent information-theoretic considerations rather than fundamental time. The admissibility framework provides the *operational* basis: time is the direction in which facts accumulate. Thermal time provides the *algebraic* implementation: the modular flow is the unique automorphism group compatible with the state's information content.

A deep question: Is the modular flow of a quantum state ω the same as the flow generated by successive commitments? If so, thermal time and commitment time would be unified—time emerges from the structure of irreversible information flow, with Tomita-Takesaki theory providing the precise mathematical form.

This remains speculative but suggests that the admissibility framework may connect to deep results in algebraic quantum field theory.

Appendix C: Comparison with Reconstruction Programs

This appendix systematically compares the admissibility approach with existing quantum reconstruction programs, showing where their axioms become theorems in our framework.

C.1 Hardy's Five Axioms

Hardy [1] derives quantum theory from:

1. Probabilities are determined by state
2. There exists a continuous reversible transformation between any two pure states
3. Simplicity: The number of degrees of freedom is minimal for given distinguishability
4. Subspaces correspond to subsystems
5. Composite systems satisfy certain tensor rules

In our framework:

- Axiom 1 becomes Theorem 7.1 (Born rule uniqueness): probabilities are forced to be $\mathrm{Tr}(E\rho)$

- Axiom 2 follows from transitive symmetry (Lemma 5.1 condition ii)
- Axiom 3 is not needed; we derive rather than assume simplicity
- Axioms 4-5 correspond to compositional locality (R4)

What we avoid assuming: Hardy assumes probabilistic structure and simplicity; we derive both.

C.2 Chiribella-D'Ariano-Perinotti (CDP)

CDP [2] derives quantum theory from:

1. Causality: No signaling from future to past
2. Perfect distinguishability: Some states can be perfectly distinguished
3. Ideal compression: Information can be maximally compressed
4. Local distinguishability: Global states distinguishable via local measurements
5. Purification: Every mixed state has a pure extension; purifications are unique up to local unitaries

In our framework:

- Axiom 1 (causality) follows from no-signaling (Proposition 8.2)
- Axiom 2 (perfect distinguishability) is part of finite distinguishability structure
- Axiom 3 (ideal compression) follows from the equivalence of operationally indistinguishable states
- Axiom 4 (local distinguishability) is local tomography, used in excluding non-complex theories
- Axiom 5 (purification) is not assumed but is compatible; it can be derived from the resulting Hilbert space structure

What we avoid assuming: CDP's key axiom is purification, which is powerful but not obviously physical. We derive Hilbert space structure without it.

C.3 Masanes-Müller

Masanes-Müller [3] derive quantum theory from:

1. Continuous reversibility: Any pair of pure states connected by continuous reversible transformation
2. Tomographic locality: Joint states determined by local measurements
3. Existence of entanglement: Some states are entangled
4. All measurements are allowed: No superselection rules

In our framework:

- Axiom 1 follows from continuous transitive symmetry
- Axiom 2 (tomographic locality) is our local tomography, used in Lemma 5.1

- Axiom 3 (entanglement existence) follows from tensor product structure under compositional locality
- Axiom 4 is related to self-duality (measurement completeness)

Distinctive feature of our approach: We begin with *facts* rather than *measurements*. Other programs assume measurement structure and derive the state space. We assume facts exist and derive both measurement and state structure.

C.4 Summary Table

Axiom/Assumption	Hardy	CDP	M-M	This work
Probabilistic structure	Assumed	Assumed	Assumed	Derived (Th 7.1)
Hilbert space	Derived	Derived	Derived	Derived (Lem 5.1)
Unitarity	Assumed	Derived	Derived	Derived (Th 5.1)
Born rule	Assumed	Derived	Derived	Derived (Th 7.1)
Tensor products	Assumed	Assumed	Assumed	Assumed (R4)
Local tomography	—	Assumed	Assumed	Used in Lem 5.1
Purification	—	Assumed	—	Derived
Facts exist	Implicit	Implicit	Implicit	**Explicit (A1)**
Finite resources	Implicit	Implicit	Implicit	**Explicit (A2)**

The distinctive contribution of this work is making the operational preconditions for physics *explicit* and deriving probabilistic structure rather than assuming it.

Appendix D: Open Problems

D.1 The Dimensionality Problem

The derivation of Dirac structure (Section 9) assumes d = 3 spatial dimensions. This is empirical input, not derived from admissibility. Can admissibility constrain dimensionality?

What's known:

- Clifford algebra representations have dimension 2^{⌊d/2⌋} for d spatial dimensions
- Bott periodicity implies representation structure repeats mod 8
- Only d = 3 gives 4-component spinors matching observed fermion structure

Possible approaches:

1. *Anthropic:* Only d = 3 supports stable structures (atoms, chemistry)

2. *Holographic:* Finite distinguishability + Bekenstein bound may constrain d
3. *Algebraic:* Complex structure forcing may select d = 3 (speculative)

Status: This remains the primary open problem. A derivation of d = 3 from admissibility would be a major result.

D.2 Bekenstein Bound and Holographic Connection

Finite distinguishability implies bounded information content. The Bekenstein bound [22] states that the maximum entropy of a region is $S \leq 2\pi RE/(\hbar c)$, where R is the radius and E is the energy. For a spherical region at fixed temperature, this becomes $S \leq A/(4\ell_P^2)$ where A is the surface area and ℓ_P is the Planck length.

Conjecture: The finite distinguishability bound N is the Bekenstein bound: $\log N \leq A/(4\ell_P^2)$.

Evidence:

- Both are operational bounds on accessible information
- Both scale with system size, not volume (holographic scaling)
- Both arise from the interface between quantum mechanics and gravity

Implications if true: Admissibility would connect directly to holography and quantum gravity. The finite distinguishability constraint wouldn't be a separate input but would follow from spacetime geometry.

D.3 Interacting Quantum Field Theory

The paper addresses kinematic QFT structure. Extending to interacting theories requires addressing:

Compatibility with Wightman axioms:

- Wightman's axioms assume: Hilbert space, Poincaré covariance, spectral condition, locality, completeness
- Admissibility provides: Hilbert space (Lemma 5.1), unitarity (Theorem 5.1), no-signaling (Proposition 8.2)
- Needed: Show Poincaré covariance follows from relativistic admissibility; derive spectral condition from energy positivity

Locality and clustering:

- Admissibility's no-signaling is compatible with Wightman locality
- Clustering (factorization of correlations at spacelike separation) should follow from compositional locality

Renormalization:

- Finite distinguishability suggests natural UV cutoff
- Question: Does admissibility select finite theories or provide a principle for renormalization?

D.4 Interaction Selection

Admissibility determines kinematic structure but not specific interactions. What additional principles might select:

- The Standard Model gauge group SU(3) × SU(2) × U(1)?
- Three generations of fermions?
- The specific coupling constants?

The VERSF program proposes that information-geometric constraints on entropy and distinguishability in curved spacetime may constrain interactions. This remains highly speculative.

Appendix E: Derived Closure Lemmas

This appendix shows that the operational postulates R1, R4, and R5 can be derived from more primitive assumptions about experimental capabilities. This strengthens the paper's claims by reducing the number of independent postulates.

E.1 Convexity from Classical Control (Derives R1)

Assumption CC (Classical Control Register). There exists an admissible two-state record/control system C with stable distinguishable states {|0⟩_C, |1⟩_C} such that:

1. C can be prepared with prescribed frequencies λ, $1-\lambda$ over repeated trials (via a classical randomizer or reproducible mixing protocol), and
2. C can conditionally route an admissible preparation procedure (if C=0 perform P_0; if C=1 perform P_1).

This is not an additional axiom about nature; it is a finite-procedure capability already implicit in admissibility. Experiments can be classically controlled and recorded—this is what "stable records exist" means operationally.

Lemma E.1 (Convexity from Classical Control). Let P_0, P_1 be admissible preparation procedures yielding operational states $\rho_0, \rho_1 \in \mathcal{S}$. Under Assumption CC, for any $\lambda \in [0,1]$ there exists an admissible preparation P_λ whose operational statistics satisfy

$p(E|\rho\_\lambda) = \lambda\, p(E|\rho_1) + (1-\lambda)\, p(E|\rho_0)$

for all admissible effects E. Therefore ρ_λ is operationally the convex mixture $\rho\_\lambda = \lambda\rho_1 + (1-\lambda)\rho_0$.

Proof. Implement P_λ by preparing the classical control register C such that over trials C=1 occurs with frequency λ and C=0 with frequency $1-\lambda$. Condition on C: if C=1 execute P_1, else execute P_0. For any admissible measurement effect E, the law of total probability for operational frequencies gives

$p(E|P\_\lambda) = p(C=1)p(E|P_1) + p(C=0)p(E|P_0) = \lambda p(E|\rho_1) + (1-\lambda)p(E|\rho_0)$.

By the definition of operational state (equivalence class of preparations agreeing on all admissible effects), P_λ represents precisely the convex mixture. ■

Comment. Convexity is no longer a postulate about the state space. It is a theorem from the existence of recordable classical control, which is already entailed by "finite procedures and records exist."

E.2 Purification from Reversible Embeddability (Derives R5)

The key insight is that irreversibility in a system S arises because degrees of freedom are discarded/inaccessible—it is not a fundamental violation of the reversible sector, only a restriction of access.

Assumption RE (Reversible Embeddability / No Fundamental Sinks). For any admissible process Φ on a system S, there exists an extended system SE and an admissible reversible transformation U on SE such that for all operational states ρ_S,

$\Phi(\rho\_S) = Tr\_E[U(\rho\_S \otimes \eta\_E)U^\dagger]$,

for some fixed operational state η_E of the environment E.

Interpretation. Irreversibility in S arises because degrees of freedom are discarded/inaccessible; it is not a fundamental violation of reversibility. This is exactly the physical content of "commitment = many-to-one because of discarded degrees."

Proposition E.2 (Reversible Sector Is Deterministic). No admissible transformation that is genuinely stochastic on operational states can be reversible (i.e., admit an admissible inverse on operational states).

Proof. Suppose a process T maps the same input operational state ρ to two distinct output operational states $\sigma_1 \neq \sigma_2$ with nonzero frequencies. Any candidate inverse T^{-1} would have to map both σ_1 and σ_2 back to ρ. Then T^{-1} is necessarily many-to-one on operational states, hence irreversibly committing distinctions, contradicting the definition of the reversible sector. ■

Theorem E.3 (Purification from Reversible Embeddability). Assume the reversible sector is unitary (or more generally, reversible and distinguishability-preserving) and Assumption RE holds. Then every mixed operational state ρ_S arises as the marginal of a pure state $|\Psi\rangle_{SE}$ of a larger system:

$$\rho_S = \mathrm{Tr}_E(|\Psi\rangle\langle\Psi|).$$

Proof. Take Φ to be the identity channel on S. By Assumption RE, there exists U on SE and fixed η_E such that $\rho_S = \mathrm{Tr}_E[U(\rho_S \otimes \eta_E)U^\dagger]$. Choose E large enough that η_E admits a purification $|\varphi\rangle_{EE'}$ *(standard in Hilbert space; alternatively, any mixed classical record state can be purified by including its classical memory register). Define* $|\Psi\rangle_{S(EE')} \equiv (U \otimes I_{E'})(|\psi\rangle_S \otimes |\varphi\rangle_{EE'})$ for any purification $|\psi\rangle_S$ of ρ_S. Tracing out EE' returns ρ_S. ■

Lemma E.4 (Uniqueness of Minimal Purification). In the Hilbert-space representation, any two minimal Stinespring dilations of the same channel are related by a unitary on the environment; hence any two purifications of ρ_S are related by a unitary on the purifying subsystem.

Proof. Standard Stinespring dilation uniqueness up to partial isometry/unitary on the ancillary space. ■

Comment. Instead of postulating purification (R5), we derive it from:

- Reversible sector cannot be stochastic (Proposition E.2)
- Irreversibility is discarding degrees of freedom (Assumption RE) ⇒ Mixedness must be marginals of pure extended states (Theorem E.3)

E.3 Tomographic Locality from Universal Controllability (Derives R4)

Assumption UC (Universal Composite Controllability). For a composite AB, the set of admissible experiments on AB is generated (to arbitrary operational precision) by:

1. Local admissible operations and measurements on A and B,
2. Shared classical control and classical communication (finite rounds), and
3. A finite universal set of admissible interactions (entangling gates) between A and B.

This is the statement that the laboratory has universal operational access to the degrees of freedom it claims to describe.

Theorem E.5 (Operational Tomography from Controllability). Under Assumption UC, if two composite states ω_{AB} and ω'_{AB} yield identical statistics for all local measurements on A and B and all their correlations (i.e., for all product effects $E_A \otimes F_B$), then they yield identical statistics for all admissible experiments on AB. Hence they are operationally identical:

$$\forall\, E_A, F_B: p(E_A, F_B \mid \omega_{AB}) = p(E_A, F_B \mid \omega'_{AB}) \Rightarrow \omega_{AB} \equiv \omega'_{AB}.$$

Proof. By Assumption UC, any admissible experiment on AB can be approximated by a circuit built from local operations, local measurements, classical control/communication, and a finite universal entangling gate set. The outcome probabilities of any such circuit are multilinear functionals of the input state evaluated on effects generated from product effects by the circuit's reversible dynamics. If two states agree on all product effects, then by linearity/continuity of operational probabilities under admissible composition (already required for reproducible finite procedures), they agree on the entire generated effect set, and therefore on all admissible experiments. By the definition of operational state as equivalence class under indistinguishability by admissible experiments, the states are identical. ■

Comment. We no longer say "Reality must be tomographically local." We say: "If your operational theory claims to describe all experimentally accessible degrees of freedom on composites (UC), then tomography follows." This makes quaternionic/exceptional "failures" read as: they contain degrees of freedom not controllable/observable under UC, hence are operationally surplus.

E.4 Operational Non-Contextuality from Effect Identity (Tightens Born Rule)

Definition (Operational Effect Identity). Two outcome-events (M, i) and (M', i') represent the same operational effect E iff for all operational states ρ,

$p(i|M, \rho) = p(i'|M', \rho)$.

An effect is an equivalence class of outcome-events with identical statistics on all states.

Lemma E.6 (Operational Non-Contextuality Is Automatic). If E is defined as an operational equivalence class of outcome-events, then the probability $p(E|\rho)$ depends only on E and ρ, not on the measurement context in which E is embedded.

Proof. If the probability depended on context, i.e., $p(i|M, \rho) \neq p(i|M', \rho)$ for some ρ, then (M, i) and (M', i) would not belong to the same equivalence class and therefore would not define the same operational effect E. Hence, for a fixed operational effect E, the probability assignment is by construction independent of context. ■

Comment. Non-contextuality is not assumed as a principle about nature; it follows from defining "same effect" in the only way compatible with operational equivalence. The Born rule derivation (Theorem 7.1) then proceeds:

1. Effects are context-independent by operational identity (Lemma E.6)
2. Additivity + continuity + Gleason/Busch-Caves → trace rule uniquely

Ontological contextuality (in the sense of Spekkens [20]) remains possible: the underlying ontic states might differ between contexts, but this cannot affect operational statistics.

E.5 Summary: Upgraded Logical Status

With the derivations in this appendix, the logical structure becomes:

Primitive assumptions:

- (A1) Operational recordability
- (A2) Finite operational resources
- (A3) Physical realizability of measurements
- (CC) Classical control register exists
- (RE) Reversible embeddability / no fundamental sinks
- (UC) Universal composite controllability

Derived constraints:

- (C1) Finite distinguishability ← Theorem 3.2
- (C2) Irreversible commitment ← Definition + A1
- (R1) Convexity ← Lemma E.1 from CC
- (R2) Continuous reversible dynamics ← forced by admissibility
- (R3) Distinguishability preservation ← forced by admissibility
- (R4) Tomographic locality ← Theorem E.5 from UC
- (R5) Purification ← Theorem E.3 from RE

Derived structure:

- Hilbert space ← Lemma 5.1 from R1–R5
- Unitarity ← Theorem 5.1
- Born rule ← Theorem 7.1 + Lemma E.6

The assumptions CC, RE, and UC are not quantum-specific axioms but minimal operational completeness requirements for a theory intended to describe all experimentally accessible degrees of freedom. They express capabilities that any laboratory doing physics must have: classical control, access to extended systems, and universal operations on composites.

References (continued)

[25] M. F. Pusey, J. Barrett, and T. Rudolph, "On the reality of the quantum state," Nature Physics **8**, 475 (2012).

[26] A. Connes and C. Rovelli, "Von Neumann algebra automorphisms and time-thermodynamics relation in generally covariant quantum theories," Class. Quantum Grav. **11**, 2899 (1994).