# Physical Necessity of Quantum Probability Structure

## A Companion to *Born Rule as Entropic Unfolding* and *The Double Square Rule*

**Keith Taylor** VERSF Theoretical Physics Program, AIDA Institute

**Primary companion paper**: This paper serves as the assumption-justification companion to *The Double Square Rule: A Derivation of Quantum Probability from Discrete Informational Geometry* (Taylor, 2025), which derives $P = |\psi|^2$ from pairwise kernel uniqueness given a set of structural axioms. The present paper shows those axioms are forced by physical admissibility. It also complements *Born Rule as Entropic Unfolding* (Taylor, 2025), which extends the geometric core to real measurement apparatus with thermodynamic corrections. Together, the three papers form a complete, non-circular derivation chain for the Born probability rule from pre-quantum physical principles.

---

## Abstract

Recent work has derived the Born probability rule $P = |\psi|^2$ as the unique stable probability assignment compatible with interference, positivity, and factorization, given structural principles including phase structure, pairwise correlation selection, equal path weighting, and compositional factorization. A common objection is that these principles are themselves assumed rather than derived. This companion paper addresses that objection directly. We show that each structural principle is **forced by physical admissibility** once one accepts three minimal facts about the world: (i) distinguishability is finite, (ii) time advances only through irreversible commitment, and (iii) composite systems must admit scalable independence. Alternatives are shown either to collapse to classical probability, violate compositional stability, or destroy temporal coherence. We prove explicit impossibility theorems for higher-order selection kernels, derive equal path weighting from geometric dependence, anchor phase normalization to minimal distinguishability, and demonstrate that non-quadratic probability rules violate normalization preservation under reversible evolution. Phase structure from continuous holonomy is shown to follow from temporal extensibility, with the core impossibility theorem proven in full and the continuity upgrade outlined as a structured proof roadmap (Appendix C); a dedicated paper will provide the complete treatment. The result is not a mathematical inevitability theorem, but a physics-level necessity result: **quantum probability structure is the only admissible organization of distinguishability consistent with time, interference, and composition**.

---

# Table of Contents

**Appendices**



---

# 1. Introduction: Scope and Purpose

## 1.1 What This Paper Does

The Born rule has been derived many times in the literature, always conditionally. Our previous papers — *Born Rule as Entropic Unfolding* (Part I) and *The Double Square Rule* (Part II) — derive quantum probability from physical mechanisms: irreversible selection acting on correlation structures between reversible paths, combined with thermodynamic refinements at the measurement boundary.

A skeptical reader may object that several structural ingredients of these derivations are assumed rather than derived:

1. Phase structure emerges from holonomy of the isometry group, but why must a discrete metric space support continuous holonomy at all?
2. Pairwise selection is privileged over individual-path or higher-order selection, but could trilinear or higher-order kernels work?
3. Path amplitudes are summed with equal weights, but shouldn't paths carry action-dependent weights as in Feynman's formulation?
4. Compositional factorization (Axiom A9) appears to encode tensor-product structure, potentially smuggling Hilbert space.
5. The fundamental representation ($n = 1$) is chosen for the phase, but what fixes the normalization?

This paper addresses each objection with explicit arguments: theorems where possible, physical necessity arguments where theorems are not yet available, and honest labeling of what remains conditional.

### 1.2 What This Paper Does Not Do

We do not claim to derive quantum mechanics from pure mathematics. Physics must begin with physical principles, and the question is never "what is logically possible?" but "what is physically admissible?" We do not produce a formal uniqueness theorem from abstract axioms. We show that within a physically motivated ontology — finite distinguishability, irreversible commitment, and compositional admissibility — **no alternative structure remains viable**.

### 1.3 Standard of Argument

Throughout this paper, we use the term **admissible** to mean:

*Capable of existing stably, being measured, and composing consistently in a universe with finite distinguishability and irreversible time advancement.*

A structure may be mathematically well-defined yet physically inadmissible because it cannot be measured, cannot compose, cannot persist, or cannot advance time. Our arguments show that the structural assumptions of the Born-rule derivation are the unique admissible options — alternatives either collapse to classical probability or produce physical pathologies.

---

## 2. Why Phase Structure Is Physically Necessary

### 2.1 The Classical Alternative

Consider a universe whose reversible distinguishability-preserving transformations form a finite permutation group $G_{fin}$. Such a system admits path counting, additive probabilities, and compositional structure. It is mathematically consistent. It is classical probability theory.

Such a universe cannot produce interference phenomena. Any theory that accounts for observed interference must admit richer reversible structure than finite permutations.

### 2.2 The Necessity Argument: Temporal Extensibility Forces Continuous Holonomy

The gap between "discrete metric space" and "continuous phase structure" is bridged by the following physical requirement:

**Temporal extensibility**: Reversible dynamics must allow arbitrarily long consistent sequences of transitions. There is no maximum path length.

This is not optional. If reversible sequences had a maximum length, the universe would reach a state from which no further reversible evolution is possible — time would stop for reversible processes.

**Theorem 2.1** (Continuous Holonomy from Temporal Extensibility):

Let ($\mathcal{S}$, d) be a finite distinguishability space (at any operational scale $\varepsilon > 0$, the number of $\varepsilon$-distinguishable states is finite). Suppose:

(A1) Time advances only through irreversible commitment of distinguishability, at least one bit per tick (TPB).

(A2) Committed distinctions persist — they remain operationally accessible (bit conservation).

(A3) Distinguishability is relational — no intrinsic labels readable without interaction.

(A4) Persistent records are accessible via loop transport (holonomy-accessibility).

(A5) There is no finite maximum number of ticks (temporal extensibility).

Then the holonomy set H cannot be finite. Furthermore, under two additional regularity assumptions:

(A6) Incremental boundedness — one reversible step changes holonomy by at most $\delta\_step$.

(A7) Total boundedness of operational balls — derived from finite distinguishability at scale.

$\bar{H}$ contains a continuous one-parameter subgroup. The minimal connected compact Abelian such subgroup is U(1).

**Proof summary** (full treatment in Appendix C):

The argument proceeds in two independent layers.

*Layer 1 — Core impossibility (from A1–A5)*: If H is finite with $|H| = K$, then only K distinct loop-effects exist. By the Holonomy Completeness Lemma (Appendix C, Section C.4 Step 2), under relational distinguishability (A3) and holonomy-accessibility (A4), distinct commitment histories can only be maintained as physically real if they map to distinct holonomy classes: $|\mathcal{C}_n\text{^\{eff\}}| \leq K$. But TPB (A1) and temporal extensibility (A5) force $|\mathcal{C}_n\text{^\{eff\}}| \geq n + 1$ for arbitrarily large n. For $n \geq K$ this is a contradiction. Therefore $|H| = \infty$. ■

*Layer 2 — Continuity upgrade (adding A6–A7)*: Incremental boundedness (A6) confines holonomies of length-n loops to a d_op-ball of radius $n \cdot \delta\_step$. Total boundedness (A7) — which follows from BCB's finite distinguishability at scale via an $\varepsilon$-net of probe protocols argument — ensures that infinitely many distinct holonomies within this ball must have Cauchy sequences and therefore accumulation points. Under the additional regularity that $\bar{H}$ is locally compact (established from A7 + metric completeness; see Appendix C, Section C.8), non-discreteness plus elements arbitrarily close to identity implies a continuous one-parameter subgroup (Gleason–Montgomery–Zippin). ■

**Status**: The core impossibility theorem (Layer 1) is a complete proof from A1–A5. The continuity upgrade (Layer 2) is structurally complete with four technical items remaining for full journal-standard rigor, identified explicitly in Appendix C, Section C.8. A dedicated paper (*Why Finite Distinguishability Forces Continuous Phase*) will provide the complete treatment of Layer 2. All other links in the Born-rule derivation chain are now explicitly justified or proven in this paper.

The impossibility of purely finite holonomy (Appendix C.4) is sufficient for the present Born-rule derivation; the continuity upgrade (Appendix C.5–C.6) is required only to identify the minimal phase group as U(1) rather than a generic infinite discrete group.

## 2.3 Why U(1) and Not a Larger Group

Given continuous holonomy, why is U(1) selected rather than SU(2), U(1) × U(1), or a larger group?

The answer is minimality and commutativity:

**Commutativity**: By Axiom A8 (relabeling invariance), the order in which path contributions are combined cannot affect probabilities. Non-Abelian phase groups (e.g., SU(2)) violate this: the order of group elements matters. Quaternionic amplitudes, which carry SU(2) phase structure, fail relabeling invariance because e^{iθ} e^{jφ} ≠ e^{jφ} e^{iθ}.

**Minimality**: U(1) is the smallest connected compact Abelian group. Larger Abelian groups (U(1) × U(1), etc.) introduce independent phase dimensions not determined by the one-dimensional holonomy of the distinguishability geometry, violating geometric dependence (Axiom A6).

**Completeness**: ℂ, the representation field of U(1), is algebraically closed, ensuring that all eigenvalue problems have solutions within the field. ℝ is not algebraically closed; extensions beyond ℂ (quaternions, octonions) sacrifice commutativity or associativity.

## 2.4 Summary

Phase structure is not an assumption — it is the minimal structure required for interference in a temporally extensible universe. A theory without continuous phase is not a different quantum theory; it is classical.

---

# 3. Why Selection Cannot Act on Individual Paths

This result is established in Part II (Double Square Rule, Section 2, Theorem: Impossibility of Individual-Path Selection) and summarized here for completeness.

## 3.1 The Impossibility Theorem

**Theorem 3.1** (from Part II): Let $P(A) = \Sigma_{\{P \in R_A\}} f(\theta(P))$ for some function $f : \mathbb{R}/2\pi\mathbb{Z} \to \mathbb{R}_{\geq 0}$. Then P(A) violates at least one of:

(i) Gauge invariance (Axiom A8): Under $\theta(P) \to \theta(P) + \alpha$, invariance forces f to be constant, eliminating all phase dependence.

(ii) Factorization (Axiom A9): The only gauge-invariant phase-dependent form $P(A) = |\psi_A|^p$ fails normalization preservation for $p \neq 2$.

(iii) Interference: Constant f yields $P(A) = c \cdot |R_A|$, which is classical path-counting with no destructive interference.

## 3.2 Physical Content

Individual-path selection assigns probabilities based on properties of single alternatives. But distinguishability is relational: $d(s_i, s_j)$ relates pairs of states. Any selection mechanism that respects the relational structure of distinguishability must depend on relations between paths, not on paths in isolation.

The impossibility theorem confirms this: individual-path rules either lose all phase information (becoming classical) or fail compositional consistency.

---

# 4. Why Higher-Order Selection Is Physically Forbidden

This is the section requiring the most substantial new content. We prove that selection kernels of order k > 2 are physically inadmissible.

## 4.1 Setup

A k-linear selection kernel assigns probability via:

$$P(A) = \Sigma_{\{P_1, \ldots, P_k \in R_A\}} W(P_1, P_2, \ldots, P_k)$$

where W depends on the phases $\theta(P_1), \ldots, \theta(P_k)$.

For k = 2, this is the bilinear kernel of the Double Square Rule. We now show k > 2 fails.

## 4.2 Impossibility of Trilinear Kernels

**Theorem 4.1** (Higher-Order Kernels Violate TPB-Consistent Normalization):

Any irreducibly k-linear selection functional with $k > 2$ either reduces to a product of bilinear functionals or has a normalization functional that is not preserved under reversible evolution (hence is physically inadmissible under TPB-consistency).

Before proving this, we establish the key mathematical fact:

**Lemma (No Universal Unitary Invariant from Mixed Symmetric Powers)**: Let $c \in \mathbb{C}^d$ with $\|c\|_2 = 1$. Any functional $Z(c)$ built as a finite linear combination of monomials involving coefficients from different tensor-power representations — i.e., terms mixing $c^{(n)}$ (n-th symmetric power) for different values of n — cannot be invariant under all U(d) transformations, unless it reduces to a function of $\|c\|_2^2$ alone.

**Proof of Lemma**: The polynomial invariants of U(d) acting on $\mathbb{C}^d$ are generated by the single invariant $c^\dagger c = \|c\|_2^2$ (this follows from Schur's lemma: the only U(d)-equivariant map $\mathbb{C}^d \to \mathbb{C}^d$ is a scalar multiple of the identity, so the only polynomial invariant of degree (p,q) in $(c, c^*)$ is proportional to $(c^\dagger c)^k$ when $p = q = k$, and vanishes otherwise). A monomial involving the n-th symmetric power $c^{(n)}$ for $n \neq \pm 1$ transforms under a representation inequivalent to the fundamental representation. Products of such monomials across different n-values transform under tensor products of inequivalent representations, whose traces are not unitary invariants. The only polynomial functional invariant under all U(d) is therefore a function of $\|c\|_2^2 = \Sigma_i |c_i|^2$.
■

**Proof of Theorem 4.1**:

Let $W : R\_A^k \to \mathbb{C}$ be a k-linear kernel satisfying geometric dependence (W depends only on phase differences) and symmetry (W is permutation-invariant). We show that for $k > 2$, the resulting probability functional either reduces to a bilinear one or violates TPB-consistent normalization.

**Step 1: General form of geometric k-linear kernels.**

By geometric dependence, W can depend only on phase differences. For $k = 3$, the independent phase differences among three paths are $\Delta_{12} = \theta(P_1) - \theta(P_2)$ and $\Delta_{13} = \theta(P_1) - \theta(P_3)$. By gauge invariance and symmetry, the most general form is:

$$W(P_1, P_2, P_3) = \Sigma_{m,n}\, a_{mn} \cdot e^{i(m\cdot\Delta_{12} + n\cdot\Delta_{13})} + \text{permutations}$$

subject to the constraint that $P(A) \in \mathbb{R}$.

**Step 2: Harmonic decomposition reveals the obstruction.**

Summing over $R_A$, any k-linear kernel with geometric dependence decomposes into products of harmonic amplitudes $\psi_A^{(n)} = \Sigma_{P \in R_A} e^{in\theta(P)}$:

$P(A) = \Sigma_{\mathbf{n}} a_{\mathbf{n}} \cdot \prod_j \psi_A^{(n_j)}$

For example, the simplest nontrivial non-cyclic trilinear kernel:

$W(P_1, P_2, P_3) = e^{i(\theta(P_1) - \theta(P_2))} \cdot e^{i(\theta(P_1) - \theta(P_3))}$

gives $P(A) = \psi_A^{(2)} \cdot (\psi_A^*)^2$, mixing the second harmonic $\psi^{(2)}$ with the fundamental conjugate $\psi^{(-1)}$.

Note that formal factorization on product systems can hold for such expressions: since $\theta(P\otimes Q) = \theta(P) + \theta(Q)$, each individual harmonic factors as $\psi_{A\otimes B}^{(n)} = \psi_A^{(n)} \cdot \psi_B^{(n)}$, and therefore $P_{XY}(A \otimes B) = P_X(A) \cdot P_Y(B)$. **Factorization is not the obstruction.**

**Step 3: The real obstruction — normalization is not reversibly invariant.**

The normalization functional is $Z(\psi) := \Sigma_A P(A)$. For physical admissibility under TPB-consistency (Definition 8.2), $Z(\psi)$ must be invariant under all reversible transformations (unitaries).

Under a unitary U, the fundamental amplitude transforms as $\psi \to U\psi$, but the n-th harmonic $\psi^{(n)}$ transforms under the n-th symmetric power representation of U — a different, inequivalent representation for $|n| > 1$. The normalization sum $Z = \Sigma_A \prod_j \psi_A^{(n_j)}$ therefore involves traces over products of distinct representation spaces.

By the Mixed Symmetric Powers Lemma above, such a functional cannot be invariant under all $U(d)$ unless it reduces to a function of $\|\psi\|_2^2$ alone — which requires all harmonic orders to be $\pm 1$ (i.e., the kernel is effectively bilinear).

For the explicit trilinear example: $Z = \Sigma_A \psi_A^{(2)} \cdot (\psi_A^*)^2$. Let $c = (1, 0)$ in $d = 2$, so $Z = 1$. Apply the Hadamard: $c' = (1/\sqrt{2})(1, 1)$. Then $\psi^{(2)} = (1/\sqrt{2})^2 + (1/\sqrt{2})^2 = 1$, $\psi^* = (1/\sqrt{2} + 1/\sqrt{2}) = \sqrt{2}$, so $Z = 1 \cdot 2 = 2 \neq 1$. Normalization is not preserved. ■

**Step 4: Positivity failure under coarse-graining.**

Even setting aside normalization, trilinear probability functionals fail positivity. Consider three paths with phases $\theta_1 = 0$, $\theta_2 = 2\pi/3$, $\theta_3 = 4\pi/3$. The non-cyclic trilinear kernel gives:

$\psi_A^{(2)} = e^0 + e^{4i\pi/3} + e^{8i\pi/3} = 1 + e^{4i\pi/3} + e^{2i\pi/3} = 0$

So $P(A) = 0$, but this vanishing is fragile: small perturbations of phases generically produce $P(A)$ with imaginary components, requiring a real-part projection that can yield negative values. Specifically, for $\theta_3 = 4\pi/3 + \varepsilon$:

$$\mathrm{Re}[P(A)] = \mathrm{Re}[\psi_A^{(2)} \cdot (\psi_A^*)^2] < 0$$

for a range of small ε, violating positivity. ■

### 4.3 General k-Linear Impossibility

**Theorem 4.2** (General Higher-Order Impossibility):

For any $k > 2$, an irreducibly k-linear probability kernel either:

(a) reduces to a product of bilinear kernels (all harmonic orders ±1), or

(b) involves higher harmonic orders ($|n| > 1$), in which case the normalization functional $Z(\psi) = \Sigma_A P(A)$ is not invariant under reversible evolution, violating TPB-consistency.

Positivity failure provides an independent secondary obstruction.

**Proof**:

A k-linear kernel with geometric dependence expands into sums of products of generalized harmonic amplitudes $\psi_A^{(n)} = \Sigma_P e^{in\theta(P)}$, with various harmonic orders n determined by the kernel structure.

Note that each individual harmonic *does* factor on product systems: $\psi_{A\otimes B}^{(n)} = \psi_A^{(n)} \cdot \psi_B^{(n)}$, since $\theta(P\otimes Q) = \theta(P) + \theta(Q)$. So the obstruction is not factorization — it is that **mixed-harmonic probability functionals fail to preserve normalization under reversible evolution**.

Any irreducibly k-linear kernel with $k > 2$ must involve harmonic orders $|n| > 1$ (otherwise it decomposes into products of bilinear $n = \pm1$ terms, giving outcome (a)). Under a unitary transformation U, the fundamental amplitude $\psi = \Sigma_P e^{i\theta(P)}$ transforms as $\psi \to U\psi$, but the n-th harmonic $\psi^{(n)} = \Sigma_P e^{in\theta(P)}$ transforms under the n-th symmetric power representation of U, not under U itself.

The normalization functional for a k-linear kernel takes the form:

$$Z(\psi) = \Sigma_A \Sigma_{\mathbf{n}} a_{\mathbf{n}} \cdot \prod_j \psi_A^{(n_j)}$$

where the harmonic indices $\{n_j\}$ include $|n_j| > 1$ for irreducibly k-linear kernels. By the Mixed Symmetric Powers Lemma (§4.2), $Z(\psi)$ cannot be invariant under all $U(d)$ unless it reduces to a function of $\|\psi\|_2^2$ alone — which requires all harmonic orders to be ±1 (i.e., the kernel is effectively bilinear).

Therefore, irreducibly k-linear kernels ($k > 2$) either:

(a) reduce to bilinear kernels (all harmonics ±1), or

(b) involve higher harmonics, in which case normalization $Z(\psi) = \Sigma_A P(A)$ is not preserved under reversible evolution — the probability "leaks" or "inflates" between measurements, violating TPB-consistency (Definition 8.2, Theorem 8.3).

The positivity argument provides a secondary, independent obstruction: mixed-harmonic functionals generically produce complex-valued or negative P(A) under small phase perturbations (see Step 4 of Theorem 4.1). ■

### 4.4 Connection to Sorkin Hierarchy

The above results are consistent with Sorkin's measure-theoretic framework, which classifies probability theories by the order of path interference. Standard quantum mechanics exhibits second-order (pairwise) interference. Third-order and higher interference has been tested experimentally and found to be absent to high precision (Sinha et al. 2010, Söllner et al. 2012). Our Theorems 4.1–4.2 provide a *derivational* explanation for this empirical fact: higher-order interference is absent because higher-order kernels are physically inadmissible.

### 4.5 Summary

Pairwise correlation is the highest-order selection structure compatible with TPB-consistent normalization. Higher-order kernels ($k > 2$) either reduce to products of bilinear terms or involve mixed harmonic orders whose normalization is not preserved under reversible evolution. Positivity failure provides an independent secondary obstruction. This is not an aesthetic choice — it is an admissibility constraint proven by explicit construction and anchored in the representation theory of U(d).

---

# 5. Why Equal Path Weights Are Forced

### 5.1 The Objection

In the Born-rule derivation, the amplitude for outcome A is:

$$\psi_A = \Sigma_{P \in R_A} e^{i\theta(P)}$$

Each path contributes with unit weight. But in Feynman's path integral, paths are weighted by $e^{iS[P]/\hbar}$ where S[P] is the classical action, which varies across paths. Why are the weights equal?

### 5.2 Weight Absorption Lemma

**Lemma 5.1** (Equal Weights from Geometric Dependence):

Let $\psi_A = \Sigma_{P \in R_A} w(P) \cdot e^{i\theta(P)}$ with $w(P) > 0$ for all P. If the correlation kernel $W(P, P')$ depends only on the phase difference $\theta(P) - \theta(P')$ (geometric dependence, Axiom A6), then $w(P) = \text{const}$.

**Proof**:

The bilinear probability is:

$$P(A) = \Sigma_{P, P' \in R_A} W(P, P')$$

The kernel has the rank-one form (Theorem 5.2 of Part II):

$$W(P, P') = \varphi(P) \cdot \varphi(P')^*$$

If paths carry non-uniform weights, the most general form compatible with the kernel structure is:

$$\varphi(P) = w(P) \cdot e^{i\theta(P)}$$

and the kernel becomes:

$$W(P, P') = w(P) \cdot w(P') \cdot e^{i(\theta(P) - \theta(P'))}$$

But Axiom A6 requires $W(P, P')$ to depend only on the geometric relationship between P and P′ — specifically, on $\theta(P) - \theta(P')$. The factor $w(P) \cdot w(P')$ depends on P and P′ individually, not only through their phase difference.

For $W(P, P')$ to depend solely on $\theta(P) - \theta(P')$, we need:

$$w(P) \cdot w(P') = \text{const for all } P, P'$$

This is satisfied if and only if $w(P) = c$ for all P (constant weight). The constant c is absorbed into the overall normalization.

**Remark on scope of geometric dependence**: Any additional positive weight $w(P)$ is itself a distinguishability label — it represents physically meaningful information about the path. By Axiom A6, all physically meaningful path information must be encoded in the distinguishability geometry (i.e., in $\theta(P)$ or in path-class multiplicity). Treating $w(P)$ as an independent scalar outside the geometry violates A6. Therefore, either $w(P)$ is absorbed into the phase definition (becoming part of θ), or it is a constant determined by symmetry. ■

## 5.3 Relation to Feynman Path Integral

In Feynman's formulation, the path integral amplitude is:

$$\psi = \int \mathcal{D}[x] \cdot e^{iS[x]/\hbar}$$

The "weight" $e^{iS/\hbar}$ is not a positive real number multiplying a separate phase — it *is* the phase. The classical action determines $\theta(P) = S[P]/\hbar$, and the integration measure $\mathcal{D}[x]$ (the "equal weighting") is the flat measure on path space, invariant under the symmetries of the theory.

Our construction is consistent with this: the phase $\theta(P)$ in our framework absorbs all path-dependent information, including what Feynman would call the "action." The equal weighting refers to the path measure, not the total contribution. Lemma 5.1 shows that this is not a choice but a consequence of geometric dependence.

### 5.4 Summary

Equal path weighting is the unique assignment compatible with geometric dependence of the correlation kernel. Non-uniform weights either violate Axiom A6 or can be absorbed into the phase definition. This is not an assumption — it is a derived constraint.

---

# 6. Why Phase Normalization Is Unique (n = 1)

### 6.1 The Objection

The U(1) group admits representations $e^{in\theta}$ for all integers n. The correlation kernel could in principle use any of these:

$W_n(P, P') = e^{in(\theta(P) - \theta(P'))}$

Why is $n = 1$ selected?

### 6.2 Faithfulness Argument

**Lemma 6.1** (Fundamental Representation from Distinguishability Preservation):

Let $\theta(P)$ be the geometric phase defined by the holonomy of the isometry group. The kernel $W_n(P, P') = e^{in(\theta(P) - \theta(P'))}$ with $|n| > 1$ fails to preserve distinguishability of paths.

**Proof**:

Two paths P, P′ are geometrically distinguishable if $\theta(P) - \theta(P') \neq 0 \pmod{2\pi}$.

For $|n| > 1$, the kernel $W_n$ identifies paths whose phase difference satisfies:

$\theta(P) - \theta(P') = 2\pi k/n$, for $k = 1, \ldots, n-1$

with paths having zero phase difference. This is because $e^{in \cdot 2\pi k/n} = e^{2\pi ik} = 1$.

These paths are physically distinguishable (they have nonzero holonomy difference) but the kernel treats them as identical. This violates the requirement that the probability rule respect the full distinguishability structure of the geometry (Axiom A6). ■

### 6.3 Anchoring to Minimal Distinguishability

The BCB framework posits a minimum distinguishable information difference $\delta_{min} > 0$. The fundamental phase generator must satisfy:

The smallest nontrivial loop produces a physically distinguishable effect.

The fundamental representation $n = 1$ maps the minimal nontrivial loop (with phase $\delta_{min}$) to $e^{i\delta_{min}} \neq 1$, which is distinguishable.

For $n > 1$, phases smaller than $2\pi/n$ are mapped closer to identity. If $\delta_{min} < 2\pi/n$, then loops with phase in $(\delta_{min}, 2\pi/n)$ are physically distinguishable but mapped to $e^{in\theta} \approx 1$ — indistinguishable from the trivial loop. This destroys information that the physics requires to be preserved.

Therefore:

The fundamental representation $n = 1$ is uniquely selected by the requirement that all physically distinguishable phases remain distinguishable in the kernel.

### 6.4 The n = −1 Equivalence

The kernel $W_{-1}(P, P') = e^{-i(\theta(P) - \theta(P'))} = W_1(P', P) = W_1(P, P')$. *Since $P(A) = \Sigma_{P,P'} W(P, P')$ and $W_{-1}$ gives $P(A) = |\psi_A|^2$* $= |\psi_A|^2$, the choice $n = -1$ produces identical probabilities. This corresponds to the complex conjugation freedom and is physically equivalent.

### 6.5 Summary

Phase normalization $n = 1$ is not conventional — it is required to preserve the full distinguishability structure of the path geometry. Higher representations lose information about physically distinct paths.

---

## 7. Why Factorization Is Physically Unavoidable

### 7.1 The Objection

Axiom A9 (factorization for independent systems) is sometimes accused of encoding tensor-product structure, which is itself a distinctively quantum axiom. Does invoking factorization amount to assuming Hilbert space?

### 7.2 Factorization as Pre-Quantum Physics

Factorization encodes something more basic than quantum mechanics:

**Independent systems must admit independent probabilities.**

This principle predates quantum theory and is present in all of classical probability, statistical mechanics, and information theory. To deny factorization is to claim that *every* system in the universe is correlated with every other system in a way that prevents any local probabilistic description.

### 7.3 Consequences of Factorization Failure

If $P(A \otimes B) \neq P(A) \cdot P(B)$ for independent systems:

1. **No local physics**: Probabilities for system X depend on what measurements are performed on distant system Y, even without interaction.
2. **No scalability**: Predictions for N-body systems require tracking all N-body correlations simultaneously, with no simplification for non-interacting subsystems.
3. **No thermodynamics**: Extensivity of entropy requires $S(A \otimes B) = S(A) + S(B)$ for independent systems, which requires probability factorization.
4. **No experimental science**: If outcomes of one experiment depend on the state of all other systems in the universe, no experiment is reproducible.

### 7.4 Factorization Is Not Hilbert-Space-Specific

Factorization is satisfied by:

- Classical probability (product measures)
- Quantum mechanics (tensor products)
- Generalized probabilistic theories (product state spaces)

It is a *structural requirement on any multi-system theory*, not a quantum axiom. The fact that quantum mechanics satisfies it via tensor products is a consequence, not an assumption.

### 7.5 Summary

Factorization is the mathematical expression of "independent systems behave independently." Any universe containing more than one system requires it. It encodes no quantum-specific structure.

---

# 8. Why Non-Quadratic Probability Rules Fail Temporally

## 8.1 The Tick-Per-Bit Constraint

The TPB principle requires:

**Time advances only through irreversible commitment of distinguishability.**

Each "tick" of time corresponds to an irreversible event that creates a definite record, consuming distinguishability resources. A probability rule must be compatible with this: the total probability must be preserved under reversible evolution (between ticks) and must produce consistent irreversible outcomes (at each tick).

## 8.2 Normalization Preservation Under Reversible Evolution

**Theorem 8.1** (Quadratic Uniqueness from Normalization Preservation):

Let $P(A) = \|\psi_A\|^p$ for some $p > 0$, where $\psi_A$ evolves unitarily: $\psi_A(t) = U(t)\cdot\psi_A(0)$ with $U^\dagger U = I$. Then $\Sigma_A P(A)$ is preserved under evolution if and only if $p = 2$.

**Proof**:

Unitary evolution preserves the inner product: $\langle U\psi|U\varphi\rangle = \langle\psi|\varphi\rangle$. In particular, it preserves $\|\psi\|^2 = \langle\psi|\psi\rangle$.

For a complete set of outcomes $\{A_i\}$ forming an orthonormal basis, $\Sigma_i |\langle A_i|\psi\rangle|^2 = \|\psi\|^2 = 1$ by Parseval's theorem.

For $p \neq 2$, consider $\Sigma_i |\langle A_i|\psi\rangle|^p$. This is the $\ell^p$ norm of the coefficient vector $\{c_i\}$ raised to the p-th power. This quantity is **not** preserved by unitary transformations for $p \neq 2$.

**Explicit example**:

Let $|\psi\rangle = |0\rangle$, so $c_0 = 1$, $c_1 = 0$.

$N_p = \Sigma_i |c_i|^p = 1^p + 0^p = 1$

Apply the Hadamard gate: $|\psi'\rangle = (1/\sqrt{2})(|0\rangle + |1\rangle)$, so $c_0' = c_1' = 1/\sqrt{2}$.

$N_p = 2 \cdot (1/\sqrt{2})^p = 2^{1 - p/2}$

For $p = 2$: $N_p = 2^{1-1} = 1$. Preserved. ✓

For $p = 4$: $N_p = 2^{1-2} = 1/2$. **Not preserved.** ✗

For $p = 1$: $N_p = 2^{1-1/2} = \sqrt{2} \approx 1.414$. **Not preserved.** ✗

Therefore, only $p = 2$ maintains $\Sigma_A P(A) = 1$ under all unitary transformations. ■

## 8.3 Formal TPB-Consistency Condition

The TPB framework gives Theorem 8.1 a precise physical interpretation through the following definition:

**Definition 8.2** (TPB-Consistency): A probability rule $P(A) = f(\psi_A)$ is **TPB-consistent** if its normalization functional $Z(\psi) := \Sigma_A f(\psi_A)$ is invariant under all reversible (isometric) transformations of $\psi$.

**Physical meaning**: Between irreversible ticks, the system evolves reversibly. TPB requires that distinguishability changes *only* at ticks — never between them. The normalization functional $Z(\psi)$ measures the total probability budget. If Z changes under reversible evolution, then distinguishability is being created or destroyed without an irreversible commitment, violating TPB.

**Theorem 8.3** (Unique TPB-Consistent Probability): The only TPB-consistent probability rule of the form $P(A) = \|\psi_A\|^p$ is $p = 2$, giving $Z(\psi) = \|\psi\|^2 = \text{const}$.

**Proof**: This is Theorem 8.1 restated in TPB language. The normalization functional $Z_p(\psi) = \Sigma_i |c_i|^p = N_p$ is invariant under all unitaries if and only if $p = 2$ (Theorem 8.1, Appendix B). Therefore $p = 2$ is the unique TPB-consistent exponent. ■

**Consequence**: TPB-consistency is not a vague interpretive principle — it is the concrete requirement that $Z(\psi)$ be a unitary invariant. The only such invariant of the form $\Sigma|c_i|^p$ is the 2-norm. This directly selects the Born rule.

## 8.4 Residual Coherence and Entropy Production

For $p = 4$, the "probability" after measurement is:

$P_i = |c_i|^4 / \Sigma_j |c_j|^4$

The renormalization factor $\Sigma_j |c_j|^4 = \text{Tr}(\rho^2)$ is the purity of the state. This varies under unitary evolution for mixed states and under basis changes for pure states, meaning:

- The entropy production per tick is basis-dependent.
- The irreversible commitment does not produce a consistent "cost" per outcome.
- Different measurement bases for the same state would require different entropy budgets, violating the requirement that the physics is determined by the state and apparatus, not by arbitrary basis choices.

In TPB language: the renormalization factor $Z_4(\psi) = \Sigma|c_i|^4$ is not a unitary invariant, so $p = 4$ fails TPB-consistency. The system's probability budget would fluctuate between ticks, creating or destroying distinguishability without irreversible commitment.

### 8.5 Summary

The Born rule (p = 2) is the unique probability exponent compatible with normalization preservation under reversible evolution. Non-quadratic rules cause probability leakage or inflation between irreversible commitments, violating temporal consistency.

---

# 9. Synthesis: The Admissibility Filter

### 9.1 Elimination Table

The following table summarizes how each alternative to the Born-rule assumptions is eliminated:

| Alternative | Failure mode | Section |
|---|---|---|
| No phase structure (discrete holonomy only) | Classical probability, no interference | §2 |
| Individual-path selection | Gauge violation or no interference | §3 |
| Trilinear selection kernel | Normalization not invariant under reversible evolution (TPB violation) | §4.2 |
| General k-linear kernel (k > 2) | Normalization failure under reversible evolution (with positivity as a secondary obstruction) | §4.3 |
| Non-uniform path weights | Violates geometric dependence (A6) | §5 |
| Higher harmonic phase (n > 1) | Distinguishability loss | §6 |
| Non-factorizing probability | No scalable multi-system physics | §7 |
| Non-quadratic exponent (p ≠ 2) | Normalization leakage under reversible evolution | §8 |

### 9.2 The Unique Survivor

Every alternative to the structural assumptions of the Born-rule derivation is eliminated by physical admissibility constraints. The unique surviving structure is:

- Continuous U(1) phase from holonomy of the isometry group
- Bilinear (pairwise) selection kernel
- Equal path weights
- Fundamental representation (n = 1)
- Factorization on product systems
- Quadratic probability (p = 2)

Together, these force:

$P(A) = |\psi_A|^2 = |\Sigma_{P \in R_A} e^{i\theta(P)}|^2$

### 9.3 The Nature of the Result

This is a **physical necessity** result, not a mathematical inevitability theorem. The distinction matters:

- A *mathematical inevitability* theorem would derive quantum probability from pure logic. No such theorem is possible in physics.
- A *physical necessity* result shows that, given minimal physical requirements (finite distinguishability, interference, composition, irreversible time), only one probability structure is admissible.

The assumptions — finite distinguishability, interference, composition, irreversible time — are empirically grounded. We do not prove they must hold; we show that if they hold, the Born rule is the unique consequence.

---

## 10. Relation to the Primary Derivation Papers

### 10.1 How This Companion Fits

The three papers in the Born-rule program serve complementary roles:

| Paper | Role | Core method |
|---|---|---|
| Part I: Entropic Unfolding | Thermodynamic refinement | MaxCal + entropy costs |
| Part II: Double Square Rule | Geometric derivation | Pairwise kernel uniqueness |
| This Companion | Assumption justification | Admissibility elimination |

Part II derives the Born rule *given* the axioms. This companion shows the axioms are *forced* by physical admissibility. Part I extends the geometric core to real apparatus with thermodynamic corrections.

### 10.2 The Complete Derivation Chain

The full non-circular derivation proceeds in four layers:

**Layer 1 — Phase emergence** (this paper, §2 + Appendix C; dedicated paper in preparation): TPB + BCB relational distinguishability + temporal extensibility ⟹ H must be infinite (Theorem, Appendix C, C.4) ⟹ H must be continuous (Corollary, Appendix C, C.5) ⟹ minimal admissible holonomy is U(1) (Proposition, Appendix C, C.6)

**Layer 2 — Structural constraints** (this paper, §§3–8): Relabeling invariance + factorization ⟹ pairwise selection (§4) Geometric dependence ⟹ equal path weights (§5) Distinguishability preservation ⟹ fundamental representation n = 1 (§6) TPB-consistency ⟹ quadratic exponent p = 2 (§8)

**Layer 3 — Born rule** (Part II, Double Square Rule): Pairwise bilinear kernel + positivity + factorization ⟹ rank-one kernel ⟹ $P(A) = |\psi_A|^2$

**Layer 4 — Thermodynamic refinement** (Part I, Entropic Unfolding): Real measurement apparatus with entropy costs ⟹ $P_i \propto |c_i|^2 \cdot e^{-\lambda\Delta S_i}$, recovering exact Born rule in the iso-entropic limit.

The Born rule is therefore not "assumed because interference exists." It is **forced because time cannot advance indefinitely without infinite holonomy (Layer 1), continuous U(1) holonomy forces phase structure which forces pairwise correlation selection under admissibility constraints (Layer 2), and pairwise selection with positivity and factorization uniquely determines quadratic probability (Layer 3)**.

### 10.3 What Remains

The core impossibility theorem — that purely finite holonomy is incompatible with TPB + BCB + temporal extensibility — is now a complete proof (Appendix C, Section C.4). The continuity upgrade from infinite to continuous holonomy (Appendix C, Section C.5) and the identification with U(1) (Appendix C, Section C.6) are structurally complete, with four technical items identified in Appendix C, Section C.8 requiring full formalization for journal-standard rigor. A dedicated paper (*Why Finite Distinguishability Forces Continuous Phase*) will provide this complete treatment. All other links in the derivation chain — pairwise selection, equal weights, fundamental representation, quadratic form — are explicitly proven in §§3–8 of this paper.

---

## 11. Discussion

### 11.1 Comparison to Other Reconstruction Programs

Several reconstruction programs derive quantum theory from operational or information-theoretic axioms:

| Program | Starting assumptions | Our comparison |
|---|---|---|
| Gleason (1957) | Hilbert space, dim ≥ 3 | We derive Hilbert space |
| Hardy (2001) | Operational primitives | We start pre-operationally |
| Chiribella et al. (2011) | Operational-probabilistic theory | Comparable minimality |
| Masanes & Müller (2011) | Information-theoretic axioms | Similar spirit, different starting point |
| Zurek (2005) | Entangled states, envariance | Requires quantum structure |

Our approach is distinctive in two respects:

1. We work *prior to Hilbert space*, deriving both the probability rule and the Hilbert space structure from informational geometry.
2. This companion paper justifies the starting axioms themselves by physical admissibility, closing the "why these axioms?" objection that all reconstruction programs face.

### 11.2 The Limits of Physical Necessity

Physical necessity arguments cannot achieve the certainty of mathematical proof. Our results hold within the physical framework of finite distinguishability and irreversible time. If these features of reality are themselves emergent from something deeper, the derivation chain would need to be extended.

We regard this as a feature, not a bug. Physics progresses by showing what follows from what, with each layer of explanation resting on empirically grounded principles. The Born rule now rests on principles that are simpler, more general, and more directly connected to observable physics than the rule itself.

---

## 12. Conclusion

The structural assumptions underlying the Born-rule derivation — phase structure, pairwise selection, equal path weighting, compositional factorization, and quadratic form — are not arbitrary mathematical choices. Each is forced by physical admissibility:

- Without continuous phase, there is no interference — only classical probability.
- Without pairwise selection, probability cannot encode relational distinguishability.
- With higher-order selection, factorization and positivity fail.
- Without equal weights, geometric dependence is violated.
- With non-fundamental phase normalization, distinguishability is destroyed.
- Without factorization, no scalable multi-system physics is possible.
- With non-quadratic exponent, normalization leaks under reversible evolution.

The Born rule does not emerge because quantum mechanics demands it. Quantum mechanics exists because the Born rule is the only admissible way to assign probabilities in a universe with finite distinguishability, irreversible time, and compositional structure.

---

# Appendix A: Technical Details for Theorem 4.1

## A.1 Harmonic Analysis of k-Linear Kernels

A general k-linear kernel with geometric dependence can be expanded in the Fourier basis of phase differences. For k paths, there are k−1 independent phase differences. The kernel admits the expansion:

$$W(P_1, \ldots, P_k) = \Sigma_{\mathbf{n}} a_{\mathbf{n}} \cdot \prod_{j=1}^{k-1} e^{i \cdot n_j \cdot (\theta(P_j) - \theta(P_k))}$$

where $\mathbf{n} = (n_1, \ldots, n_{k-1}) \in \mathbb{Z}^{k-1}$.

The probability becomes:

$$P(A) = \Sigma_{\mathbf{n}} a_{\mathbf{n}} \cdot \prod_{j=1}^{k-1} \psi_A^{(n_j)} \cdot (\psi_A^{(-1)})^{(\text{from } P_k)}$$

where $\psi_A^{(n)} = \Sigma_P e^{in\theta(P)}$.

## A.2 Factorization and Normalization Constraints on Harmonics

For product systems, each individual harmonic factors correctly: $\psi_{A \otimes B}^{(n)} = \Sigma_{P,Q} e^{in(\theta(P) + \theta(Q))} = [\Sigma_P e^{in\theta(P)}] \cdot [\Sigma_Q e^{in\theta(Q)}] = \psi_A^{(n)} \cdot \psi_B^{(n)}$. This holds for all n.

The obstruction for $k > 2$ is therefore not that individual harmonics fail to factor, but that **normalization of mixed-harmonic probability functionals is not preserved under unitary evolution**. The probability functional:

$$P(A) = \Sigma_{\mathbf{n}} a_{\mathbf{n}} \prod_j \psi_A^{(n_j)}$$

involves products of amplitudes at different harmonic orders. Under a unitary transformation U, the fundamental amplitude $\psi^{(1)}$ transforms as $\psi \to U\psi$, but the n-th harmonic $\psi^{(n)}$ transforms under the n-th symmetric power of U — a different representation. The normalization sum $\Sigma_A P(A)$ therefore involves traces over products of distinct representation spaces, which are not unitary invariants unless all harmonics are ±1.

For the factorization of the probability itself (not individual harmonics), the requirement $P(A \otimes B) = P(A) \cdot P(B)$ applied to the harmonic expansion gives:

$$\Sigma_{\mathbf{n}} a_{\mathbf{n}} \prod_j \psi_{A \otimes B}^{(n_j)} = [\Sigma_{\mathbf{n}} a_{\mathbf{n}} \prod_j \psi_A^{(n_j)}] \cdot [\Sigma_{\mathbf{n}} a_{\mathbf{n}} \prod_j \psi_B^{(n_j)}]$$

Since each $\psi^{(n_j)}$ factors individually, the left side becomes $\Sigma_{\mathbf{n}} a_{\mathbf{n}} \prod_j \psi_A^{(n_j)} \cdot \psi_B^{(n_j)}$. For this to equal the product on the right for arbitrary $\psi_A$, $\psi_B$, the sum must contain only a single term (rank-one in harmonic space). The surviving term must satisfy $\Sigma_j n_j = 0$ (gauge invariance). For $k = 2$, this gives $n_1 = -n_2 = \pm 1$ (the bilinear kernel). For $k > 2$, no solution with $|n_j| \leq 1$ and $\Sigma\, n_j = 0$ generates an irreducibly k-linear kernel — all such solutions factor into products of bilinear terms. Solutions with $|n_j| > 1$ formally factorize but violate normalization preservation, as shown in Section 4.3 and Theorem 8.1.

---

# Appendix B: Normalization Preservation Proof (General Case)

**Theorem B.1**: Let $\{c_i\}_{i=1}^d$ with $\Sigma\, |c_i|^2 = 1$ be the coefficients of a state in an orthonormal basis. Define $N_p := \Sigma_i\, |c_i|^p$ for $p > 0$. Then $N_p$ is invariant under all unitary transformations $c_i \to \Sigma_j\, U_{ij}\, c_j$ if and only if $p = 2$.

**Proof**:

($p = 2$): $N_2 = \Sigma\, |c_i|^2 = \|\mathbf{c}\|^2$, which is preserved by unitarity. ✓

($p \neq 2$): It suffices to exhibit one unitary transformation that changes $N_p$.

Let $d = 2$, $\mathbf{c} = (1, 0)$, so $N_p = 1$. Apply the Hadamard transformation:

$\mathbf{c}' = (1/\sqrt{2})(1, 1)$

Then $N_p = 2 \cdot (1/\sqrt{2})^p = 2^{1 - p/2}$.

For $p = 2$: $N_p = 1$. ✓ For $p \neq 2$: $N_p = 2^{1-p/2} \neq 1$. ✗

Since a single counterexample suffices, $N_p$ is not unitarily invariant for $p \neq 2$. ■

---

# Appendix C: TPB → Continuous Holonomy — Full Proof Roadmap

This appendix provides the complete proof outline for Theorem 2.1, designed to serve simultaneously as a self-contained roadmap for the dedicated paper (*Why Finite Distinguishability Forces Continuous Phase*).

The argument proceeds in two layers. The first layer is a complete combinatorial proof that purely finite holonomy is impossible under TPB + BCB (Section C.4). The second layer upgrades "infinite" to "continuous" using two additional physical assumptions (Section C.5). The two layers are logically independent — the impossibility theorem stands on its own.

### C.1 Goal and Claim

**Goal**: Show that a universe with finite relational distinguishability and Tick-Per-Bit (TPB) time advancement cannot have purely finite reversible holonomy; it must admit a continuous one-parameter holonomy subgroup, with the minimal admissible connected compact Abelian case being U(1).

**Informal claim**: If time can advance without bound through irreversible commitments, and if committed information must remain operationally accessible through relational distinguishability, then the reversible structure must support unboundedly many distinct loop effects. Purely finite holonomy saturates after finitely many loop compositions, making accumulated commitments operationally indistinguishable and violating bit conservation. Therefore, admissible reversible structure requires an infinite — and, given incremental boundedness, continuous — holonomy parameter.

## C.2 Assumptions

We state seven assumptions. A1–A5 are used in the core impossibility theorem (Section C.4). A6–A7 are used only for the continuity upgrade (Section C.5).

**A1. TPB (Tick-Per-Bit)**: Each tick commits at least one genuinely new distinguishability commitment. In particular, after n ticks, the set of physically distinct commitment histories has size $|\mathcal{C}_n$^{eff}$| \geq n + 1$. (At minimum: "0 ticks happened" vs "1 tick happened" vs … vs "n ticks happened" must remain distinguishable.) We use the weakest possible growth rate — linear — sufficient for the contradiction argument. Stronger bounds (exponential, etc.) would strengthen the result but are not needed.

**A2. Bit conservation / persistence**: Committed distinctions do not evaporate. Formally, if two commitment histories $\sigma \neq \sigma'$ are distinct at time n, there exists some later time and some operational procedure that can still distinguish them. Otherwise earlier "commitments" were not real.

**A3. Relational distinguishability (no intrinsic labels)**: There is no absolute readout of history by "looking at a label." Distinguishability is established only by operations that compare states and structures. Operationally: to extract a record, one must run a probe through the structure and compare its returned state to a reference. This is the non-negotiable "no direct-access RAM labels" premise.

**A4. Holonomy-accessibility of records**: Any persistent record that is operationally accessible through reversible operations is accessible via the action of loops (transport-and-return) on probes. Equivalently: the observable downstream effect of a record on reversible probes is captured by loop holonomy classes.

Two clarifications on the scope of A4:

*Between-ticks restriction*: Between ticks, all operations are reversible (by the reversible micro-dynamics assumption). Any verification protocol that does not itself commit a new tick must therefore be a reversible transport-and-return — i.e., a loop. Verification via irreversible measurement is allowed, but it consumes a tick, and the *content* of what's being verified (the distinction between σ and σ′) must be encoded in something the reversible transport can detect — which is holonomy. The information accessible without spending a tick is precisely the holonomy.

*Loops need not be spatial*: A "loop" need not involve spatial transport. Any reversible interaction that begins and ends with a probe in a reference configuration — including local coupling to a system — constitutes a loop in the relevant sense. Coupling a probe to a system, letting them interact reversibly, and returning the probe to its reference state is operationally: prepare probe → transport through interaction → return probe. The "holonomy" is the state change acquired by the probe during the interaction.

**A5. Temporal extensibility**: There is no finite maximum number of ticks / no global time horizon in ordinary regimes. If the universe halted after finitely many ticks, no sustained physics — and no science — would be possible. This is a minimal admissibility condition.

**A6. Incremental boundedness** *(used only in Section C.5)*: There exists a constant $\delta_{step} > 0$ such that appending a single reversible step to a loop changes its operational holonomy by at most $\delta_{step}$ in the operational metric $d_{op}$. That is, for any loop L and any single reversible map T:

$$d_{op}(\mathrm{hol}(T \circ L), \mathrm{hol}(L)) \leq \delta_{step}$$

This is a continuity-of-effects assumption, not a geometric smoothness assumption. It says that one physical operation produces a bounded physical effect. Without it, arbitrarily small changes in procedure could produce arbitrarily large physical effects — a universe where adding one reversible step could change holonomy by an unbounded amount would be operationally pathological (no stable transport, no reproducible protocols). A6 is required for reproducibility: if one additional reversible operation could change outcomes arbitrarily, no stable experimental protocol could exist.

**A7. Total boundedness of operational balls** *(used only in Section C.5; derived from finite distinguishability at scale)*: For each $R > 0$ and each $\varepsilon > 0$, the ball $B_R(e) \subset (H, d_{op})$ can be covered by finitely many $d_{op}$-balls of radius ε. This follows from BCB's finite distinguishability at scale (Definition 1) — see Section C.5 for the derivation.

We do **not** assume Hilbert space, complex numbers, phases, or interference.

## C.3 Definitions

**(1) Distinguishability space**: A set $\mathcal{S}$ of states equipped with a distinguishability relation or metric d, sufficient to define when two states are operationally distinct. The metric is relational: d(s, s′) is established by comparison protocols, not by reading intrinsic labels. **Finite distinguishability** means: at any operational scale $\varepsilon > 0$, the number of ε-distinguishable states is finite. (This is compatible with $\mathcal{S}$ itself being countably infinite, but prevents infinite precision at any finite resolution.)

**(2) Reversible transformation group G**: The set of reversible maps $T : \mathcal{S} \to \mathcal{S}$ that preserve distinguishability (isometries or their discrete analogue). Composition gives a group. G acts on $\mathcal{S}$ by bijections satisfying $d(T(s), T(s')) = d(s, s')$.

**(3) Loops and holonomy**: Fix a reference state $s_0$. A loop of length n is a finite composition of reversible maps returning $s_0$ to itself:

$L = T_n \circ \cdots \circ T_1, L(s_0) = s_0$

The **holonomy** of a loop L is the residual internal transformation induced by L on the local distinguishability neighborhood of $s_0$ — i.e., the induced action of L on states near $s_0$, modulo the identity. Two loops with different holonomies produce different transformations of the local structure even though both return $s_0$ to itself.

**(4) Holonomy set at depth n**: Let $H_n$ be the set of distinct holonomies achievable by loops of length $\leq n$. Define the total holonomy set $H = \bigcup_n H_n$. Since loop composition gives group structure, H is a subgroup of the stabilizer of $s_0$ in G.

**(5) Operational distinguishability of holonomies**: Two holonomies $h, h' \in H$ are **operationally distinct** if there exists a probe protocol — a state s near $s_0$ and a subsequent measurement — such that transporting s under h versus h′ produces distinguishable outcomes. This defines a pseudometric on H:

d_op(h, h′) := sup over probe protocols of |observable difference when transporting under h vs h′|

Quotienting by d_op = 0 yields a metric space (H/∼, d_op). This is the **operational topology** on holonomies. A holonomy class is "trivial" if d_op(h, e) = 0 — it has no detectable effect on any probe.

**(6) Commitment history space**: After n ticks, the system has a commitment history $\sigma_n = (\sigma_1, \sigma_2, \ldots, \sigma_n)$, where each $\sigma_i$ records the outcome of the i-th irreversible commitment. The set of possible commitment histories after n ticks is $\mathcal{C}_n$. Under TPB (at least one bit per tick), $|\mathcal{C}_n| \geq n + 1$, and generically $|\mathcal{C}_n|$ grows at least linearly in n.

**(7) Purely finite holonomy**: H is **purely finite** if $|H| < \infty$ — the total number of operationally distinct holonomy classes is finite, regardless of loop length. This is the physical definition: if two holonomies cannot be operationally distinguished, they are the same in BCB. (Note: this is stronger than "discrete topology" in the group-theoretic sense, which would include infinite discrete groups like $\mathbb{Z}$. C.4 rules out $|H| < \infty$. C.5 targets the remaining case $|H| = \infty$ but H discrete — e.g., $\mathbb{Z}$ — and shows this too is excluded under incremental boundedness.)

## C.4 Core Impossibility Theorem

**Theorem (Temporal Extensibility Forbids Purely Finite Holonomy)**: Under assumptions A1–A5, the holonomy set H cannot be finite.

**Proof** (by contradiction):

**Step 1 — If H is finite, there is a finite bound on distinct verification outcomes.**

Assume H is purely finite. Then there exists a finite number K such that:

$|H| = K < \infty$

No matter what loops are run, there are only K distinct operational loop-effects available to probes. Repeated loop concatenation cannot produce new holonomy classes — the holonomy spectrum **saturates**.

**Step 2 — Persistent committed histories must be distinguishable by some loop-effect.**

**Lemma (Holonomy Completeness for Reversible Readout)**: Under A3–A4, if two commitment histories $\sigma \neq \sigma'$ are operationally distinguishable without consuming an additional tick, then there exists a loop/probe protocol whose holonomy classes differ between σ and σ′. Therefore the equivalence relation "indistinguishable by any loop" coarsens $\mathcal{C}_n$ to $\mathcal{C}_n$^{eff}, and $\mathcal{C}_n$^{eff} injects into H_eff.

**Proof of Lemma**: By A2 (persistence), any two distinct commitment histories $\sigma \neq \sigma'$ must remain distinguishable by some operational test at some later time. By A3 (relational distinguishability), there are no intrinsic labels — distinguishability requires comparison via physical operations. By A4 (holonomy-accessibility), any such comparison that does not itself consume a tick reduces to a probe transport-and-return procedure — i.e., a loop — whose returned probe state differs between σ and σ′.

Define the equivalence relation $\sigma \sim \sigma'$ iff no loop/probe protocol produces different holonomy classes for σ vs σ′. The quotient $\mathcal{C}_n/\sim$ is precisely $\mathcal{C}_n$^{eff} — the set of commitment histories that are genuinely distinguishable via reversible readout. By construction, distinct elements of $\mathcal{C}_n$^{eff} map to distinct holonomy classes, giving an injection:

$\mathcal{C}_n$^{eff} ↪ H_eff ⟹ |$\mathcal{C}_n$^{eff}| ≤ |H_eff|

Under our assumption $|H| = K < \infty$, we have |H_eff| ≤ K, so:

|$\mathcal{C}_n$^{eff}| ≤ K

If you can only ever observe K distinct loop effects, you can only ever maintain K distinct histories as physically real, because holonomy is the only operational channel by which they remain distinguishable. ■

**Step 3 — TPB forces unbounded growth of distinct histories.**

By A1 (TPB) and A5 (temporal extensibility), after n ticks:

$|\mathcal{C}_n\text{^\{eff\}}| \geq n + 1$

and n can be made arbitrarily large.

**Step 4 — Contradiction for large n.**

Combine Steps 2 and 3:

$n + 1 \leq |\mathcal{C}_n\text{^\{eff\}}| \leq K$

This is impossible for $n \geq K$. Therefore the assumption $|H| < \infty$ is false. ■

**What the theorem establishes**: The proof shows that if time can advance without bound via genuine commitments (A1 + A5), and those commitments must remain physically real (A2), and records are relationally accessible only through reversible probe transport (A3 + A4), then there must exist unboundedly many operationally distinct loop effects. Purely finite holonomy is ruled out.

**Addressing the classical computer objection**: A classical computer has trivial holonomy — all loops act as identity on internal states — yet processes unbounded information by growing its state space (adding memory registers). Why can't a TPB universe do the same?

The answer is that classical information is readable via **direct access** to memory registers. You can "just look" at a bit in RAM without transporting a probe around a loop. But this direct-access readability is itself a physical assumption: it requires that the memory state carries an intrinsic, non-relational label distinguishable without comparison.

A3 (relational distinguishability) denies this. In a BCB universe, there are no intrinsic labels. The only way to determine whether register R holds 0 or 1 is to compare R against a reference — and comparison is a transport-and-return operation, i.e., a loop with holonomy. A classical computer with direct-access memory is a universe with **trivially readable absolute labels**, which A3 excludes by construction.

The classical computer is not a counterexample — it is a system that violates A3. Within BCB, holonomy is the **unique channel** through which committed information is operationally accessible.

## C.5 From Infinite to Continuous Holonomy

The impossibility theorem (C.4) establishes $|H| = \infty$. This section upgrades "infinite" to "continuous," using the additional assumptions A6 (incremental boundedness) and A7 (total boundedness). The two results are logically independent: the impossibility theorem stands on its own; the continuity upgrade is a corollary with additional physical input.

**The gap between infinite and continuous**: An infinite group need not be continuous. The integers $\mathbb{Z}$ under addition are infinite but discrete — every element is isolated. If H were isomorphic to $\mathbb{Z}$ (or any infinite discrete group), the impossibility theorem would be satisfied but continuous phase structure would not follow. We need to show that H has accumulation points in the operational topology.

**The continuity argument**: As n increases, $H_n$ contains increasingly many distinct holonomy classes (from the impossibility theorem). The operational topology (Definition 5) provides a metric d_op on H.

Given A6 (incremental boundedness), loops of length n produce holonomies within a d_op-ball of radius n · δ_step around the identity. For any fixed n, all holonomies in $H_n$ lie within this ball. As $n \to \infty$, the number of distinct holonomies within the ball grows without bound (by the impossibility theorem), but the ball radius grows only linearly.

To conclude that this forces accumulation points, we invoke A7 (total boundedness of operational balls):

**A7. Total boundedness of operational balls** *(derived from BCB finite distinguishability)*: For each R > 0 and each ε > 0, the ball B_R(e) ⊂ (H, d_op) can be covered by finitely many d_op-balls of radius ε.

This is not an independent assumption — it follows from BCB's finite distinguishability at scale. The argument proceeds in three steps. First, at any resolution ε, the set of operationally distinct probe protocols is itself finite up to equivalence: a probe protocol consists of preparing a probe state, transporting it under the holonomy, and measuring the result, and both the probe states and measurements are drawn from a distinguishability space where only finitely many ε-distinct elements exist at any finite scale (Definition 1). Second, d_op is defined as a supremum over probe protocols (Definition 5), but at resolution ε this supremum can be replaced by a maximum over a finite ε-net of protocols — the remaining protocols contribute at most ε additional distinguishing power. Third, within any bounded effect region (a d_op-ball of radius R), two holonomies can be ε-separated only if they produce different outcomes on at least one protocol in this finite ε-net. Since each protocol has finitely many ε-distinguishable outcomes, the number of holonomies that can be mutually ε-separated within B_R(e) is finite. This is precisely total boundedness.

Given A6 + A7: the ball B_{nδ}(e) is totally bounded and contains $|H_n| \to \infty$ distinct elements. In a totally bounded metric space, any infinite subset has a Cauchy sequence — and therefore accumulation points in the completion. (This is the correct replacement for a Bolzano–Weierstrass argument, which requires compactness or sequential compactness that we have not yet established.)

Therefore the closure $\bar{H}$ in the operational topology is not discrete — it contains limit points.

Additionally, compositions of nearly-canceling paths ($L \circ L^{-1}$ with small perturbation — i.e., a loop followed by its near-inverse) produce holonomies arbitrarily close to identity. This provides the "elements arbitrarily close to identity" needed for the following standard result:

Under the additional regularity that $\bar{H}$ is a locally compact topological group, non-discreteness plus elements arbitrarily close to the identity implies the existence of a continuous one-parameter subgroup. (This is a standard structure theorem for locally compact groups, often referred to as the Gleason–Montgomery–Zippin theorem; see C.8(i)–(ii) for what must be verified rigorously.) We show $\bar{H}$ is locally compact by proving that operational balls are totally bounded (A7) and complete under d_op — a totally bounded complete metric space is compact, so closed balls in $\bar{H}$ are compact, establishing local compactness.

**Conclusion**: $\bar{H}$ contains a continuous one-parameter subgroup. The holonomy spectrum is not merely infinite — it is continuous.

**Corollary (Continuous Holonomy)**: Under assumptions A1–A7, the closure of the holonomy set $\bar{H}$ in the operational topology contains a continuous one-parameter subgroup.

## C.6 Minimal Continuous Holonomy Is U(1)

**Claim**: The minimal continuous holonomy subgroup compatible with admissibility constraints is U(1).

**Argument in three parts**:

**(6a) Abelian (order-neutral).** Relabeling invariance (Axiom A8 from the Born-rule derivation) requires that the combined effect of independent subloops in the same equivalence class does not depend on concatenation ordering. For the holonomy connected component, this means:

$h_1 \circ h_2 = h_2 \circ h_1$

for all holonomies $h_1$, $h_2$ in the connected component. Therefore the connected component is Abelian.

**(6b) One-dimensional.**

**Definition**: The holonomy connected component has **dimension d_H** if there exist d_H independent continuous parameters that can be independently read out via loop-effect protocols — i.e., d_H continuous functions $\varphi_1, \ldots, \varphi_{d_H} : H_{conn} \to \mathbb{R}$ such that each $\varphi_i$ can be independently varied by choice of loop, and each is operationally measurable.

TPB requires exactly one irreducible commitment per tick. Each tick creates one bit of commitment, which requires one dimension of holonomy to track (by the Holonomy Completeness Lemma in C.4, Step 2). If $d_H \geq 2$, a single tick would enable two independent continuous readouts from one irreducible commitment — two independently distinguishable transport effects from a single one-bit event. But an irreducible commitment, by definition, produces exactly one bit of distinguishability. Two independent holonomy dimensions would allow two independent continuous observables to change under a single irreducible commitment, contradicting TPB's "one bit per tick" irreducibility.

More precisely: the holonomy dimension $d_H$ satisfies $d_H \geq 1$ (from the continuity corollary, C.5) and $d_H \leq 1$ (from TPB irreducibility). Therefore $d_H = 1$.

**(6c) Compact.** The holonomy acts on a finite distinguishability neighborhood (BCB). Compactness follows if the holonomy acts faithfully on a finite-dimensional bounded distinguishability neighborhood: a continuous group preserving a bounded metric on a finite-dimensional space is a subgroup of a compact group. BCB's finite distinguishability at scale (Definition 1) ensures the neighborhood is finite-dimensional — at any resolution ε, only finitely many states are ε-distinguishable, so the effective dimension of the probe space is finite. Therefore the holonomy connected component is compact.

**(6d) Conclusion.** The minimal connected, compact, one-dimensional, Abelian group is:

$U(1) \cong \mathbb{R}/2\pi\mathbb{Z}$

**Proposition (Minimal Admissible Holonomy)**: If the holonomy connected component is one-dimensional, compact, and Abelian, it is isomorphic to U(1).

## C.7 U(1) Holonomy Yields Phase Differences and Interference

**Claim**: Once U(1) holonomy is established, phase structure and interference emerge as direct consequences.

**Argument**: With U(1) holonomy, each loop L has a well-defined holonomy angle $\theta(L) \in \mathbb{R}/2\pi\mathbb{Z}$. For two reversible paths P, P′ leading to the same macro-outcome A:

- Form the loop $L = P \circ (P')^{-1}$ (go forward along P, return along P′ reversed).
- Its holonomy is $e^{i(\theta(P) - \theta(P'))}$.
- Observable consequences depend on the phase **difference** $\theta(P) - \theta(P')$, not on absolute phases (since the absolute phase is gauge-dependent: shifting all phases by a constant α does not change any holonomy of the form $P \circ (P')^{-1}$).

This is exactly the phase structure assumed in the Born-rule derivation (Part II, Section 3). It is not put in by hand — it is a faithful representation of the emergent U(1) holonomy forced by TPB + BCB + temporal extensibility.

**Interference** follows immediately: when multiple paths contribute to the same outcome, their combined effect depends on pairwise phase differences through holonomy. Constructive combination occurs when phase differences are near 0; destructive combination occurs when they are near π. This is not assumed — it is the operational content of U(1) holonomy acting on probe states.

**Corollary (Phase Structure)**: The amplitude $\psi_A = \Sigma_{P \in R_A} e^{i\theta(P)}$ is a faithful representation of the emergent U(1) holonomy structure, not an independent assumption. The "edge phase" construction of Part II is a consequence of the holonomy forced by A1–A7.

## C.8 What Remains to Be Made Fully Rigorous

The core impossibility theorem (C.4) is a complete proof from assumptions A1–A5. The continuity upgrade (C.5) and the identification with U(1) (C.6) require additional work to reach full publication standard. Specifically:

**(i) Operational topology formalization (C.5)**: Define d_op precisely as a mathematical object. Show it is a metric (not just a pseudometric) after quotienting. Prove that A7 (total boundedness) follows from BCB finite distinguishability at scale — specifically, that the finite ε-covering property of the distinguishability space (Definition 1) implies finite ε-covering of holonomy balls in d_op.

**(ii) Nearly-canceling paths (C.5)**: Prove rigorously that compositions of the form $L \circ T\_\varepsilon \circ L^{-1}$ (where T_ε is a "small" reversible step) produce holonomies within ε · δ_step of identity. This is the key technical step connecting incremental boundedness (A6) to the accumulation-point argument. Show that the one-parameter subgroup theorem applies to ($\bar{H}$, d_op) — this requires showing $\bar{H}$ is a locally compact metrizable topological group.

**(iii) Dimensionality (C.6b)**: The argument that d_H = 1 from TPB irreducibility needs careful formalization. Prove that independent holonomy dimensions (as defined in C.6b) correspond to independent commitment channels, and that an irreducible commitment cannot produce independent effects along two holonomy directions. The key is showing that "independently readable via loop protocols" implies "independently committable," which contradicts TPB irreducibility.

**(iv) Compactness (C.6c)**: Show rigorously that BCB's finite distinguishability at scale forces the holonomy connected component to be compact. The argument requires: (a) the holonomy acts faithfully on a finite-dimensional space (from BCB's finite ε-covering), and (b) it preserves a bounded metric (from the distinguishability structure). Together these place the holonomy inside a compact group.

Items (i)–(iv) are technical completions, not structural gaps. The logical architecture of the proof is in place; what remains is making each step satisfy journal-standard rigor. The TPB → continuous holonomy result is the **sole remaining foundational proof obligation** in the complete derivation of quantum probability from pre-quantum physical principles.

## C.9 Relation to Existing Literature

The claim that holonomy structure is forced by information-theoretic constraints connects to several existing programs:

**Sorkin's quantum measure theory** classifies theories by interference order. Our derivation provides a physical *reason* for second-order interference: it is the order corresponding to pairwise holonomy, which is what one-dimensional U(1) holonomy naturally produces.

**Hardy's and Chiribella et al.'s reconstructions** assume operational primitives and derive quantum structure. Our approach works at a deeper level: we derive the operational primitives themselves (phase, interference, Hilbert space) from pre-operational physical constraints (distinguishability, time, commitment).

**Wheeler's "it from bit"** proposed that physical reality arises from information. Our derivation chain makes this precise: the Born rule (and with it, quantum mechanics) arises from the requirements of bit conservation and commitment under relational distinguishability.

**Landauer's principle** connects information erasure to thermodynamic cost. Our TPB principle is a temporal analogue: information commitment has a temporal cost (one tick per bit). The connection to Landauer is made explicit in Part I (Entropic Unfolding), where thermodynamic costs produce the entropic correction to the Born rule.

---

# References


1. Born, M. (1926). Zur Quantenmechanik der Stoßvorgänge. *Zeitschrift für Physik*, 37, 863–867.
2. Gleason, A.M. (1957). Measures on the closed subspaces of a Hilbert space. *Journal of Mathematics and Mechanics*, 6, 885–893.
3. Hardy, L. (2001). Quantum theory from five reasonable axioms. *arXiv:quant-ph/0101012*.
4. Masanes, L. & Müller, M.P. (2011). A derivation of quantum theory from physical requirements. *New Journal of Physics*, 13, 063001.
5. Chiribella, G., D'Ariano, G.M., & Perinotti, P. (2011). Informational derivation of quantum theory. *Physical Review A*, 84, 012311.
6. Zurek, W.H. (2005). Probabilities from entanglement, Born's rule from envariance. *Physical Review A*, 71, 052105.
7. Busch, P. (2003). Quantum states and generalized observables: a simple proof of Gleason's theorem. *Physical Review Letters*, 91, 120403.
8. Sorkin, R.D. (1994). Quantum mechanics as quantum measure theory. *Modern Physics Letters A*, 9, 3119–3127.
9. Sinha, U. et al. (2010). Ruling out multi-order interference in quantum mechanics. *Science*, 329, 418–421.
10. Landauer, R. (1961). Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5, 183–191.
11. Taylor, K. (2025). Born Rule as Entropic Unfolding. *VERSF Theoretical Physics Program, AIDA Institute*.
12. Taylor, K. (2025). The Double Square Rule: A Derivation of Quantum Probability from Discrete Informational Geometry. *VERSF Theoretical Physics Program, AIDA Institute*.
13. Wheeler, J.A. (1990). Information, physics, quantum: the search for links. In *Complexity, Entropy and the Physics of Information*, ed. W.H. Zurek, Addison-Wesley.
14. Montgomery, D. & Zippin, L. (1955). *Topological Transformation Groups*. Interscience Publishers.