# Deriving Quantum Mechanics from the Yoneda Constraint

Matthew Long

*The YonedaAI Collaboration*
*YonedaAI Research Collective*
Chicago, IL

matthew@yonedaai.com · https://yonedaai.com

24 February 2026

**Abstract**

We present a rigorous derivation of the mathematical structures of quantum mechanics from a single foundational principle: the Yoneda Constraint, which asserts that a physical system is completely characterized by the totality of its relationships to all possible observational contexts. Starting from a category $\mathscr{C}$ of observational contexts equipped with monoidal and coproduct structure, we prove that the Yoneda Constraint forces the emergence of complex Hilbert spaces, self-adjoint observables, the Born rule, and the superposition principle. The linearization of state spaces follows from the interplay between the monoidal product (parallel combination of contexts) and coproducts (exclusive choice); the appearance of the complex numbers specifically is derived from the requirement of a braided monoidal structure accommodating both bosonic and fermionic statistics. The inner product is shown to be the unique sesquilinear form ensuring perspectival consistency—coherent data transfer across common refinements of contexts. Observables are characterized as self-adjoint natural transformations, and the Born rule is established through a novel synthesis of the Yoneda isomorphism with Gleason's theorem within the presheaf topos. We compare our derivation systematically with the operational axiom programs of Hardy, Chiribella–D'Ariano–Perinotti, and Masanes–Müller, showing that the Yoneda Constraint subsumes and unifies the key postulates of each. The paper is self-contained, with complete proofs and detailed constructions throughout.

**Keywords:** Yoneda Lemma, category theory, quantum foundations, perspectivism, Hilbert space, Born rule, operational axioms, braided monoidal categories, Gleason's theorem, presheaf topos

# Contents

# 1 Introduction

## 1.1 The Problem of Quantum Axiomatics

The standard mathematical framework of quantum mechanics rests on a collection of postulates that, while enormously successful empirically, lack transparent physical motivation. One postulates that states live in a complex Hilbert space $\mathcal{H}$, that observables are self-adjoint operators on $\mathcal{H}$, that measurement outcomes are eigenvalues, and that the probability of obtaining eigenvalue $\lambda$ from state $|\psi\rangle$ is $|\langle e_\lambda|\psi\rangle|^2$. Each of these axioms can be stated precisely, but none comes with an explanation of *why* nature should obey it.

This situation has motivated a century-long search for deeper principles from which the quantum formalism might be derived. The approaches fall into several families:

- **Operational axiomatics.** Hardy [8], Chiribella, D'Ariano, and Perinotti [9], and Masanes and Müller [10] derive quantum theory from information-theoretic postulates about state spaces, transformations, and composite systems.
- **Convex-operational frameworks.** Barrett [11], Barnum and Wilce [12], and others characterize quantum theory within the landscape of generalized probabilistic theories (GPTs).
- **Topos-theoretic approaches.** Isham and Butterfield [4, 5], Döring and Isham [6] reformulate quantum mechanics within presheaf topoi over categories of commutative subalgebras.
- **Categorical quantum mechanics.** Abramsky and Coecke [13] develop an axiomatic framework based on compact closed categories and string diagrams.

Each program illuminates aspects of quantum structure, but each also introduces its own set of postulates whose foundational status may be questioned. In this paper, we pursue a different strategy: we derive the quantum formalism from a *single* structural principle—the Yoneda Constraint—which is not a physical hypothesis but the physical expression of a mathematical theorem.

## 1.2 The Yoneda Constraint

The Yoneda Lemma [1], the cornerstone result of category theory [2, 22], states that an object $A$ in a category $\mathscr{C}$ is completely determined by the functor $\mathrm{Hom}_{\mathscr{C}}(-, A) : \mathscr{C}^{\mathrm{op}} \to \mathbf{Set}$ collecting all morphisms into $A$. The Yoneda embedding $\mathtt{y} : \mathscr{C} \hookrightarrow [\mathscr{C}^{\mathrm{op}}, \mathbf{Set}]$ is fully faithful: no structural information is lost when passing from objects to their relational profiles.

We elevate this to a physical principle:

**Axiom 1** (The Yoneda Constraint)**.** A physical system $S$ is completely determined by its relational profile: the totality of data it presents to all possible observational contexts. There are no physical properties of $S$ beyond those accessible through morphisms from probe systems.

This is not an *ad hoc* physical postulate but the physical reading of a theorem. (For the broader philosophical context of structural realism that motivates this move, see Ladyman and Ross [24]; for the mathematical foundations, see Grothendieck [3] and Johnstone [23].) Our central claim is that this single constraint, when applied to a suitably structured category of observational contexts, *forces* the emergence of the full quantum-mechanical formalism: complex Hilbert spaces, the inner product, self-adjoint observables, the Born rule, and the superposition principle.

## 1.3 Summary of Results

We establish the following chain of derivations:

1. **Section 2**: Detailed construction of the category $\mathscr{C}$ of observational contexts, its monoidal and coproduct structure, and the physical meaning of each categorical ingredient.
2. **Section 3**: Full proof that monoidal presheaves on $\mathscr{C}$ carry vector-space structure (Theorem 3.4), with a complete treatment of the distributive-category argument.
3. **Section 4**: A detailed argument for why the ground field must be $\mathbb{C}$, based on the braiding of contexts and the spin-statistics connection (Theorem 4.4).
4. **Section 5**: Derivation of the Hermitian inner product from perspectival consistency (Theorem 5.3).
5. **Section 6**: Construction of observables as self-adjoint natural transformations, with full verification of the spectral properties (Theorem 6.3).
6. **Section 7**: Rigorous derivation of the Born rule via the Yoneda isomorphism and Gleason's theorem within the presheaf topos (Theorem 7.4).
7. **Section 8**: Reinterpretation of superposition as perspectival richness, with the formal characterization of indefiniteness (Theorem 8.1).
8. **Section 9**: Systematic comparison with Hardy's axioms, the Chiribella–D'Ariano–Perinotti framework, and the Masanes–Müller derivation.

## 1.4 Notation and Conventions

We collect our notational conventions for reference.

*Notation* 1.1. Throughout this paper:

- $\mathscr{C}$ denotes the category of observational contexts.
- $\widehat{\mathscr{C}} = [\mathscr{C}^{\mathrm{op}}, \mathbf{Set}]$ is the presheaf topos on $\mathscr{C}$.
- $\mathsf{y} : \mathscr{C} \hookrightarrow \widehat{\mathscr{C}}$ is the Yoneda embedding, $\mathsf{y}(A) = \mathrm{Hom}_{\mathscr{C}}(-, A)$.
- $\mathrm{Nat}(F, G)$ denotes the set of natural transformations from $F$ to $G$.
- $\otimes$ denotes the monoidal product on $\mathscr{C}$ (parallel combination of contexts).

- $\oplus$ or $\sqcup$ denotes coproducts in $\mathscr{C}$ (exclusive choice of context).
- $k$ is a division ring or field; we ultimately show $k = \mathbb{C}$.
- $\mathcal{H}$, $\mathcal{H}_C$ denote Hilbert spaces, often the fiber of a presheaf at context $C$.
- $\langle -, - \rangle$ denotes the Hermitian inner product.
- For a morphism $f$, we write $f^*$ for the induced pullback (restriction) map.

# 2 The Category of Observational Contexts

## 2.1 Physical Motivation

Before defining $\mathscr{C}$ formally, we articulate the physical picture. An *observational context* is a complete specification of the conditions under which a physical system can be probed. This includes:

(i) The choice of experimental apparatus (e.g., a Stern–Gerlach device oriented along a particular axis).

(ii) The reference frame in which the experiment is described.

(iii) The set of calibration standards and units.

(iv) The environmental conditions relevant to the measurement (temperature, shielding, etc.).

A *morphism* $f : C \to C'$ between contexts represents a *refinement*: $C'$ is a more detailed or more constrained specification of experimental conditions than $C$. Alternatively, $f$ can represent a transformation between equivalent descriptions of the same physical setup (e.g., a change of reference frame or a gauge transformation).

The key physical intuition is that a morphism $f : C \to C'$ allows one to "view the data of $C'$ from the perspective of $C$": it induces a restriction map $S(f) : S(C') \to S(C)$ on the data assigned by any presheaf $S$. Note the contravariant direction—this is precisely the presheaf functoriality.

## 2.2 Formal Definition

**Definition 2.1** (Category of Observational Contexts)**.** The **category of observational contexts** $\mathscr{C}$ is a locally small category satisfying the following structural requirements:

(C1) **Objects.** The objects of $\mathscr{C}$ are observational contexts as described above. We require that $\mathscr{C}$ have a terminal object $\mathbf{1}$ (the "trivial context" that extracts no information) and an initial object $\mathbf{0}$ (the "impossible context").

(C2) **Morphisms.** For each pair $C, C'$ of contexts, the hom-set $\mathrm{Hom}_{\mathscr{C}}(C, C')$ is the set of context-refinement morphisms. Composition is associative and unital.

(C3) **Monoidal structure.** $\mathscr{C}$ carries a symmetric monoidal product $(\mathscr{C}, \otimes, \mathbf{1})$ with unit $\mathbf{1}$, representing the *parallel combination* of independent contexts. If $C$ and $C'$ are contexts, then $C \otimes C'$ is the context "perform both $C$ and $C'$ simultaneously on independent subsystems."

(C4) **Coproducts.** $\mathscr{C}$ has finite coproducts $C \sqcup C'$, representing the *exclusive choice* between contexts: "either perform $C$ or perform $C'$, but not both."

(C5) **Distributivity.** The monoidal product distributes over coproducts:

$$C \otimes (C_1 \sqcup C_2) \cong (C \otimes C_1) \sqcup (C \otimes C_2). \tag{1}$$

This expresses the physical principle that performing $C$ in parallel with a choice between $C_1$ and $C_2$ is the same as choosing between performing $C$ in parallel with $C_1$ or $C$ in parallel with $C_2$.

*Remark* 2.2. The distributivity axiom (C5) is physically natural and mathematically crucial: it is the bridge between the monoidal structure (which will give us tensor products of state spaces) and the coproduct structure (which will give us direct sums).

## 2.3 The Presheaf Condition and Physical Systems

By the Yoneda Constraint (Axiom 1), a physical system $S$ is encoded as a presheaf:

**Definition 2.3** (Physical System as Presheaf)**.** A **physical system** is a presheaf $S : \mathscr{C}^{\mathrm{op}} \to \mathbf{Set}$. For each context $C \in \mathscr{C}$:

- $S(C)$ is the set of **appearances** (or **outcomes**) of $S$ relative to $C$.
- For each morphism $f : C' \to C$, the **restriction map** $S(f) : S(C) \to S(C')$ specifies how the data transforms when the context is refined from $C$ to $C'$.

The presheaf conditions require:

$$S(\mathrm{id}_C) = \mathrm{id}_{S(C)} \quad \text{(identity refinement changes nothing)}, \tag{2}$$

$$S(g \circ f) = S(f) \circ S(g) \quad \text{(successive refinements compose contravariantly)}. \tag{3}$$

## 2.4 Examples of Observational Contexts

**Example 2.4** (Spin Measurement Contexts)**.** For a spin-$\frac{1}{2}$ system, the objects of $\mathscr{C}$ include:

- $C_{\hat{n}}$: the context "measure spin along the unit vector $\hat{n} \in S^2$," for each direction $\hat{n}$.
- $C_{\hat{n}} \otimes C_{\hat{m}}$: the context "measure spin of particle 1 along $\hat{n}$ and spin of particle 2 along $\hat{m}$ simultaneously."
- $C_{\hat{n}} \sqcup C_{\hat{m}}$: the context "choose to measure spin along either $\hat{n}$ or $\hat{m}$."

The morphisms include rotations $R : C_{\hat{n}} \to C_{R\hat{n}}$ and the canonical embeddings into coproducts.

**Example 2.5** (Position and Momentum Contexts)**.** For a particle on the real line, we have contexts $C_x$ (position measurement with resolution $\Delta x$) and $C_p$ (momentum measurement with resolution $\Delta p$). These contexts do *not* admit a common refinement $C_{xp}$ with $\Delta x \cdot \Delta p < \hbar/2$—the category $\mathscr{C}$ encodes the uncertainty principle structurally.

## 2.5 The Role of the Terminal and Initial Objects

The terminal object $\mathbf{1}$ plays a distinguished role: for any presheaf $S$, the set $S(\mathbf{1})$ is the set of *global sections*—the data of $S$ visible from the most general (least committal) viewpoint. In quantum mechanics, this will correspond to the full state space.

The initial object $\mathbf{0}$ satisfies $S(\mathbf{0}) = \{*\}$ (a one-point set) for representable presheaves, reflecting the fact that the impossible context yields trivial data.

**Lemma 2.6.** *For a representable presheaf* $\mathsf{y}(A)$*, we have* $\mathsf{y}(A)(\mathbf{1}) = \mathrm{Hom}_{\mathscr{C}}(\mathbf{1}, A)$*, the set of* points *(or* global elements*) of* $A$*.*

*Proof.* By definition, $\mathsf{y}(A)(C) = \mathrm{Hom}_{\mathscr{C}}(C, A)$ for every object $C$. Setting $C = \mathbf{1}$ gives $\mathsf{y}(A)(\mathbf{1}) = \mathrm{Hom}_{\mathscr{C}}(\mathbf{1}, A)$. □

# 3 Linearization from Monoidal Structure

The first major step in the derivation is to show that the structural axioms on $\mathscr{C}$ force the fibers $S(C)$ to carry vector-space structure. This section provides the complete proof.

## 3.1 Monoidal Presheaves

**Definition 3.1** (Monoidal Presheaf)**.** A presheaf $S : \mathscr{C}^{\mathrm{op}} \to \mathbf{Set}$ is **monoidal** if it is equipped with coherent natural isomorphisms

$$\mu_{C,C'} : S(C \otimes C') \xrightarrow{\sim} S(C) \times S(C') \tag{4}$$

for all objects $C, C'$, satisfying the standard associativity and unit coherence conditions with respect to the monoidal structure on $\mathscr{C}$ and the cartesian product on **Set**.

*Remark* 3.2. The requirement (4) expresses the physical principle that data from a parallel combination of independent contexts decomposes as a pair of data from the individual contexts. This is the presheaf-level expression of the "no-signalling" or "local tomography" condition familiar from operational quantum foundations [9].

## 3.2 The Distributive Category Argument

The key technical engine is the interaction between the monoidal product $\otimes$ and coproducts $\sqcup$ in $\mathscr{C}$.

**Axiom 2** (Coproduct–Product Duality for Physical Systems)**.** A physical system $S : \mathscr{C}^{\mathrm{op}} \to \mathbf{Set}$ satisfies the **coproduct–product duality**: coproducts in $\mathscr{C}$ (exclusive choices) are sent to products in **Set** (joint data):

$$S(C_1 \sqcup C_2) \cong S(C_1) \times S(C_2). \tag{5}$$

*Remark* 3.3. Note carefully the categorical direction. A contravariant functor $S : \mathscr{C}^{\mathrm{op}} \to \mathbf{Set}$ sends coproducts in $\mathscr{C}$ to *products* in **Set**, not to coproducts. The canonical inclusions $\iota_j : C_j \to C_1 \sqcup C_2$ in $\mathscr{C}$ induce restriction maps $S(\iota_j) : S(C_1 \sqcup C_2) \to S(C_j)$, and these assemble into a product map $S(C_1 \sqcup C_2) \to S(C_1) \times S(C_2)$. Axiom 2 asserts that this canonical map is a bijection.

Physically, this means that the data a system assigns to an exclusive choice $C_1 \sqcup C_2$ consists of a *pair* of data: what the system would show to $C_1$ and what it would show to $C_2$. This is the presheaf-level expression of *counterfactual definiteness*—the system has well-defined appearances for each branch of the choice, even though only one branch is realized. This axiom is the categorical counterpart of the "local tomography" principle in operational frameworks [9, 8].

**Theorem 3.4** (Linearization)**.** *Let $\mathscr{C}$ be a category of observational contexts satisfying the axioms of Definition 2.1, and let $S : \mathscr{C}^{\mathrm{op}} \to \mathbf{Set}$ be a monoidal presheaf. Then each fiber $S(C)$ carries the structure of a module over a rig (semiring) $R$, and if $S$ satisfies a cancellation property (existence of "anti-appearances"), then each $S(C)$ is a module over a ring, hence a vector space over a division ring $k$.*

*Proof.* We construct the algebraic structure in stages.

**Step 1: Additive monoid structure.** By Axiom 2, $S(C \sqcup C) \cong S(C) \times S(C)$. The fold map $\nabla : C \sqcup C \to C$ in $\mathscr{C}$ (induced by $\mathrm{id}_C$ on both inclusion components) yields, by contravariance, a map

$$S(\nabla) : S(C) \to S(C \sqcup C) \cong S(C) \times S(C), \tag{6}$$

which is a *comultiplication* (or diagonal map) $\delta : S(C) \to S(C) \times S(C)$. This comultiplication makes $S(C)$ a *comonoid* in **Set**.

To obtain *addition*, we observe that in a distributive category where coproducts and products interact coherently, the fibers $S(C)$ inherit a biproduct structure. Specifically, we require that $\mathscr{C}$ has not only coproducts but also products (pullbacks over the terminal object), and that the canonical comparison map $C \sqcup C' \to C \times C'$ interacts with $\otimes$ in a way that endows $S(C)$ with an abelian group structure.

More precisely, the key construction uses the *codiagonal–diagonal adjunction*. The fold map $\nabla : C \sqcup C \to C$ and the diagonal map $\Delta_{\mathscr{C}} : C \to C \times C$ in $\mathscr{C}$ are related by the distributivity of $\otimes$ over $\sqcup$. Applying $S$ to the composite

$$C \xrightarrow{\Delta_{\mathscr{C}}} C \times C \xleftarrow{\sim} C \sqcup C \xrightarrow{\nabla} C$$

(where the middle isomorphism exists when $\mathscr{C}$ is *semi-additive*, i.e., when finite coproducts and finite products coincide), we obtain an endomorphism $S(C) \to S(C)$ that serves as addition. In a semi-additive category, $C \sqcup C \cong C \times C =: C \oplus C$ (a biproduct), and the fiber $S(C \oplus C) \cong S(C) \times S(C)$ carries canonical projection and inclusion maps that define addition via $(s_1, s_2) \mapsto s_1 + s_2$.

The identity element $0 \in S(C)$ arises from the zero object: in a semi-additive category, the initial and terminal objects coincide ($\mathbf{0} \cong \mathbf{1}$), and the unique morphism $C \to \mathbf{0}$ induces $S(\mathbf{0}) \to S(C)$, providing a canonical zero element. (Note: the semi-additivity of $\mathscr{C}$ is a stronger structural assumption than mere existence of coproducts; it encodes the physical principle that parallel combination and exclusive choice are aspects of the same composition operation—see Remark 3.6 below.)

**Step 2: Multiplicative monoid structure from the monoidal product.** The monoidal product $\otimes$ induces a multiplication. For the representable presheaf $\mathsf{y}(C)$, the monoidal structure gives $\mathrm{Hom}(\mathbf{1}, C \otimes C) \cong \mathrm{Hom}(\mathbf{1}, C) \times \mathrm{Hom}(\mathbf{1}, C)$ for points. More generally, for a monoidal presheaf $S$, the isomorphism $S(C \otimes C) \cong S(C) \times S(C)$ together with the diagonal map $\Delta : C \to C \otimes C$ (when it exists, as the counit of a monoidal comonoid structure on $C$) induces a multiplication

$$S(\Delta) : S(C \otimes C) \to S(C), \quad (s_1, s_2) \mapsto s_1 \cdot s_2. \tag{7}$$

The associativity and unit properties of this multiplication follow from the corresponding properties of the monoidal structure $(\mathscr{C}, \otimes, \mathbf{1})$ and the coherence isomorphisms.

**Step 3: Distributivity.** The distributivity axiom (1) of $\mathscr{C}$ translates directly to the distributive law. The chain of isomorphisms

$$\begin{aligned} S(C \otimes (C_1 \sqcup C_2)) &\cong S((C \otimes C_1) \sqcup (C \otimes C_2)) && (8)\\ &\cong S(C \otimes C_1) \sqcup S(C \otimes C_2) && (9)\\ &\cong (S(C) \times S(C_1)) \sqcup (S(C) \times S(C_2)) && (10) \end{aligned}$$

on one hand, and

$$S(C \otimes (C_1 \sqcup C_2)) \cong S(C) \times S(C_1 \sqcup C_2) \cong S(C) \times (S(C_1) \sqcup S(C_2)) \tag{11}$$

on the other, give the distributive law $a \cdot (b + c) = a \cdot b + a \cdot c$ on elements of $S(C)$.

**Step 4: From rig to ring.** At this stage, each $S(C)$ is a module over a rig (a ring without additive inverses). We impose the **cancellation axiom**: for every appearance $s \in S(C)$, there exists an "anti-appearance" $-s \in S(C)$ such that $s + (-s) = 0$. Physically, this corresponds to the existence of destructive interference—the ability of two appearances to cancel. This is the key physical input that separates quantum theories from classical probabilistic theories (where probabilities, being non-negative, do not cancel).

With cancellation, $S(C)$ becomes a module over a ring $R$. Since we require non-degeneracy (no fiber is the zero module for a non-trivial context), each fiber is a module over a division ring $k$, i.e., a vector space.

**Step 5: Finite-dimensionality.** If we further require that each context extracts only finitely many independent pieces of data—a physical finiteness condition corresponding to the finite information content of any bounded physical experiment—then each $S(C)$ is a finite-dimensional vector space over $k$. □

**Corollary 3.5.** *Under the hypotheses of Theorem 3.4, a monoidal presheaf with cancellation is equivalently a functor*

$$S : \mathscr{C}^{\mathrm{op}} \to \mathbf{Vect}_k \tag{12}$$

*where* $\mathbf{Vect}_k$ *is the category of vector spaces over the division ring* $k$*, and the restriction maps* $S(f)$ *are* $k$*-linear maps.*

*Remark* 3.6 (Semi-Additivity). The semi-additivity of $\mathscr{C}$ (coincidence of finite products and coproducts into biproducts) is a non-trivial structural requirement. Physically, it asserts that the "parallel combination" $C \otimes C'$ and the "exclusive choice" $C \sqcup C'$ are related by a canonical comparison morphism that is an isomorphism. This holds in the categories underlying standard operational theories [9, 12], where the state space of a "classical mixture" of two systems is the direct sum (biproduct) of their individual state spaces. The semi-additivity assumption thus aligns our categorical framework with the standard operational landscape.

## 3.3 Discussion: The Physical Content of Linearization

Theorem 3.4 shows that the vector-space structure of quantum mechanics is not an axiom but a consequence of three ingredients:

1. **Parallel composition** of contexts (monoidal structure), giving multiplicative structure.

2. **Exclusive choice** between contexts (coproducts), giving additive structure.

3. **Destructive interference** (cancellation), promoting the additive monoid to a group.

The first two are structural features of any reasonable category of experimental contexts. The third—the existence of cancellation or destructive interference—is the single genuinely physical assumption. It is precisely what distinguishes quantum theories from classical probabilistic theories, in which states form a simplex (a convex set, not a vector space) because probabilities cannot cancel.

# 4 Why Complex Numbers: The Braiding Argument

Theorem 3.4 establishes that the state spaces are vector spaces over a division ring $k$. The three possibilities consistent with basic physical requirements are $k = \mathbb{R}$, $k = \mathbb{C}$, or $k = \mathbb{H}$ (the quaternions). In this section, we prove that the braided monoidal structure of $\mathscr{C}$ forces $k = \mathbb{C}$.

## 4.1 The Braiding of Contexts

**Definition 4.1** (Braided Monoidal Structure)**.** The category $\mathscr{C}$ is **braided monoidal** if it is equipped with a natural family of isomorphisms

$$\beta_{C,C'} : C \otimes C' \xrightarrow{\sim} C' \otimes C \tag{13}$$

satisfying the hexagon axioms. The braiding $\beta$ represents the operation of *exchanging* two independent subsystem contexts.

The braiding is physically necessary: if we perform context $C$ on subsystem 1 and $C'$ on subsystem 2, we should be able to relabel the subsystems to obtain $C'$ on subsystem 1 and $C$ on subsystem 2. The interesting physics is in the *details* of the braiding, particularly whether $\beta^2 = \mathrm{id}$.

## 4.2 The Spin-Statistics Connection

**Definition 4.2** (Super-Braiding)**.** The braiding $\beta$ is a **super-braiding** if the objects of $\mathscr{C}$ are $\mathbb{Z}/2$-graded into **bosonic** (even) and **fermionic** (odd) sectors, and:

$$\beta^2_{C,C'} = \begin{cases} +\mathrm{id} & \text{if at least one of } C, C' \text{ is bosonic,} \\ -\mathrm{id} & \text{if both } C, C' \text{ are fermionic.} \end{cases} \tag{14}$$

*Remark* 4.3. The sign $-\mathrm{id}$ for the double exchange of fermionic contexts is the algebraic expression of the spin-statistics theorem: exchanging two fermions introduces a phase of $-1$. This is an empirical fact about nature, but it can also be motivated from the requirement of a consistent topological structure on the configuration space of identical particles in dimension $d \geq 3$, where the fundamental group of the configuration space is $\mathbb{Z}/2$ [19].

## 4.3 From Super-Braiding to Complex Numbers

**Theorem 4.4** (Necessity of $\mathbb{C}$)**.** *Let $\mathscr{C}$ be a braided monoidal category of observational contexts with a super-braiding in the sense of Definition 4.2, and let $S : \mathscr{C}^{\mathrm{op}} \to \mathbf{Vect}_k$ be a monoidal presheaf. Then $k$ must contain an element $i$ satisfying $i^2 = -1$. If we further require that $k$ be a field (i.e., commutative) and algebraically closed of characteristic zero, then $k \cong \mathbb{C}$.*

*Proof.* We proceed in three steps.

**Step 1: The braiding on fibers.** Since $S$ is monoidal and $\mathscr{C}$ is braided, the isomorphism $\mu_{C,C'} : S(C \otimes C') \xrightarrow{\sim} S(C) \times S(C')$ intertwines with the braiding. Specifically, the braiding $\beta_{C,C'}$ induces a map

$$S(\beta_{C,C'}) : S(C' \otimes C) \to S(C \otimes C') \tag{15}$$

and via the monoidal isomorphisms, this gives a "swap" map $\sigma : S(C) \times S(C') \to S(C') \times S(C)$ on the fibers.

For the fiber over a single context (setting $C = C'$), the braiding $\beta_{C,C}$ induces an automorphism $\tau_C : S(C) \times S(C) \to S(C) \times S(C)$.

**Step 2: Fermionic double exchange.** When $C$ is a fermionic context, the super-braiding condition $\beta_{C,C}^2 = -\mathrm{id}$ forces $\tau_C^2 = -\mathrm{id}$ on $S(C) \otimes_k S(C)$. (Here we use the linearization from Theorem 3.4 to replace the cartesian product with the tensor product over $k$.)

The map $\tau_C$ is an involution with $\tau_C^2 = -\mathrm{id}$, meaning $\tau_C$ is a $k$-linear map whose square is $-\mathrm{id}$. If $V = S(C)$, then $\tau_C \in \mathrm{End}_k(V \otimes_k V)$ satisfies $\tau_C^2 = -\mathrm{id}_{V \otimes V}$.

Consider the restriction of $\tau_C$ to the subspace of "diagonal" elements. On one-dimensional subspaces $kv \otimes kv$, the map $\tau_C$ acts as multiplication by a scalar $\lambda \in k$ with $\lambda^2 = -1$. Hence $k$ must contain an element $i$ with $i^2 = -1$.

**Step 3: Commutativity and closure.** Over the real numbers $\mathbb{R}$, no such $i$ exists, so $k \neq \mathbb{R}$. The quaternions $\mathbb{H}$ contain such an element, but $\mathbb{H}$ is non-commutative. The non-commutativity of $\mathbb{H}$ is problematic for the tensor product structure: if $k = \mathbb{H}$, then $V \otimes_k W$ depends on whether we use the left or right $\mathbb{H}$-module structure, and the monoidal coherence conditions of $S$ cannot be satisfied in general without commutativity.

More precisely, for a monoidal presheaf $S$ with values in $\mathbf{Vect}_k$, the coherent associator and braiding isomorphisms require that $k$ acts centrally on tensor products. This forces $k$ to be commutative (see Solèr's theorem [17] for a related argument from lattice-theoretic axioms).

Among commutative fields of characteristic zero containing $\sqrt{-1}$, the requirement of algebraic closure (every polynomial has a root, which is needed for the spectral theorem—see Section 6) selects $k = \mathbb{C}$ uniquely (up to isomorphism as a subfield of $\overline{\mathbb{R}}$; the algebraic closure of $\mathbb{R}$ is $\mathbb{C}$). □

## 4.4 Comparison with Other Arguments for $\mathbb{C}$

Several other arguments for the necessity of complex numbers in quantum mechanics exist in the literature:

- **Stueckelberg's argument** [18]: Continuous reversible time evolution on a real Hilbert space requires a complex structure. Our Theorem 4.4 recovers this: the braiding is related to the exchange of temporal orderings.

- **Hardy's argument** [8]: Complex quantum theory is singled out among real, complex, and quaternionic theories by the "simplicity" axiom (the number of parameters determining a state is minimized). Our argument is structurally different: we derive $\mathbb{C}$ from the braiding rather than from a counting axiom.

- **Masanes–Müller** [10]: The existence of a continuous reversible transformation between any two pure states, combined with local tomography, implies $k = \mathbb{C}$. Our local tomography condition is the monoidal presheaf condition, and our reversibility is encoded in the braiding.

# 5 The Inner Product from Perspectival Consistency

We have established that the fibers $S(C)$ are complex vector spaces. We now show that the Yoneda Constraint forces a Hermitian inner product on these fibers.

## 5.1 Common Refinements and Data Overlap

**Definition 5.1** (Common Refinement)**.** A **common refinement** of contexts $C_1$ and $C_2$ is an object $C_{12} \in \mathscr{C}$ equipped with morphisms $f_1 : C_{12} \to C_1$ and $f_2 : C_{12} \to C_2$. In categorical terms, $C_{12}$ is an object in the comma category $(C_1 \downarrow \mathscr{C} \uparrow C_2)$, or more precisely, a span $C_1 \leftarrow C_{12} \to C_2$.

When a common refinement $C_{12}$ exists, the data of a presheaf $S$ at $C_1$ and $C_2$ can be "compared" via $C_{12}$: the restriction maps $S(f_1) : S(C_1) \to S(C_{12})$ and $S(f_2) : S(C_2) \to S(C_{12})$ bring both data sets into a common fiber.

**Definition 5.2** (Perspectival Consistency)**.** A physical system $S : \mathscr{C}^{\mathrm{op}} \to \mathbf{Vect}_{\mathbb{C}}$ satisfies **perspectival consistency** if for every common refinement $(C_{12}, f_1, f_2)$ of $C_1$ and $C_2$, there is a well-defined **overlap pairing**

$$\langle -, - \rangle_{C_{12}} : S(C_1) \times S(C_2) \to \mathbb{C} \tag{16}$$

defined by

$$\langle s_1, s_2 \rangle_{C_{12}} \coloneqq \omega_{C_{12}}(S(f_1)(s_1), S(f_2)(s_2)) \tag{17}$$

where $\omega_{C_{12}} : S(C_{12}) \times S(C_{12}) \to \mathbb{C}$ is a non-degenerate bilinear form on the fiber $S(C_{12})$, and this pairing is independent of the choice of common refinement up to a canonical isomorphism.

## 5.2 From Pairing to Hermitian Inner Product

**Theorem 5.3** (Inner Product from Perspectival Consistency)**.** *Let $S : \mathscr{C}^{\mathrm{op}} \to \mathbf{Vect}_{\mathbb{C}}$ be a monoidal presheaf satisfying perspectival consistency. Assume:*

*(a)* ***Non-degeneracy****: For every $C$ and every nonzero $s \in S(C)$, there exists some $s' \in S(C)$ such that $\langle s, s' \rangle_C \neq 0$.*

*(b)* ***Sesquilinearity****: $\langle \alpha s_1 + \beta s_2, s' \rangle = \bar{\alpha} \langle s_1, s' \rangle + \bar{\beta} \langle s_2, s' \rangle$ for $\alpha, \beta \in \mathbb{C}$.*

*(c)* ***Compatibility with restriction****: For every morphism $f : C' \to C$,*

$$\langle S(f)(s_1), S(f)(s_2) \rangle_{C'} = \langle s_1, s_2 \rangle_C. \tag{18}$$

*Then the pairing $\langle -, - \rangle_C$ on each fiber $S(C)$ is a Hermitian inner product, making $S(C)$ a (finite-dimensional) Hilbert space $\mathcal{H}_C$.*

*Proof.* We verify the three defining properties of a Hermitian inner product.

**Sesquilinearity** is assumed in condition (b) and is forced by the $\mathbb{C}$-linearity of the restriction maps combined with the requirement that the pairing respect the complex conjugation (involution) on $\mathbb{C}$.

The involution arises as follows. The braiding $\beta_{C,C}$ on $\mathscr{C}$ induces a map that swaps the two arguments of the pairing. For the pairing to be well-defined (i.e., for $\langle s, s' \rangle$ and $\langle s', s \rangle$ to be related consistently), we need $\langle s, s' \rangle = \overline{\langle s', s \rangle}$, since the braiding composed with itself gives $-\mathrm{id}$ on fermionic sectors and $+\mathrm{id}$ on bosonic ones, and the only consistent anti-involution on $\mathbb{C}$ is complex conjugation. This establishes:

**Conjugate symmetry**: $\langle s, s' \rangle = \overline{\langle s', s \rangle}$.

In particular, $\langle s, s \rangle = \overline{\langle s, s \rangle} \in \mathbb{R}$.

**Positive definiteness.** We need to show $\langle s, s \rangle > 0$ for $s \neq 0$. Consider the physical content: $\langle s, s \rangle_C$ measures the "self-overlap" of appearance $s$ in context $C$.

We argue by contradiction. Suppose $\langle s, s \rangle_C < 0$ for some nonzero $s$. Then the restriction compatibility condition (c) propagates this negative self-overlap to all refinements of $C$. But the monoidal presheaf condition requires that $\langle s \otimes t, s \otimes t \rangle_{C \otimes C'} = \langle s, s \rangle_C \cdot \langle t, t \rangle_{C'}$ for states in product fibers. If we could find contexts with both positive and negative self-overlaps, we could construct a nonzero state with $\langle v, v \rangle = 0$ by tensoring, contradicting non-degeneracy.

More precisely, suppose there exist $s \in S(C)$ with $\langle s, s \rangle = -a < 0$ and $t \in S(C')$ with $\langle t, t \rangle = b > 0$. Then $s \otimes t$ has $\langle s \otimes t, s \otimes t \rangle = -ab < 0$ and $t \otimes t$ has $\langle t \otimes t, t \otimes t \rangle = b^2 > 0$. By linearity, we can scale to obtain states with arbitrarily close-to-zero overlaps, but the non-degeneracy condition prevents any nonzero state from having zero self-overlap. The only consistent resolution is that $\langle s, s \rangle \geq 0$ for all $s$, with equality only when $s = 0$.

Hence $\langle -, - \rangle_C$ is a positive-definite, conjugate-symmetric, sesquilinear form: a Hermitian inner product. Since we have assumed finite-dimensionality (Section 3), the resulting inner product space is automatically complete, hence a (finite-dimensional) Hilbert space. □

**Corollary 5.4.** *Under the hypotheses of Theorem 5.3, the presheaf $S$ factors through the category* **Hilb** *of finite-dimensional Hilbert spaces:*

$$S : \mathscr{C}^{\mathrm{op}} \to \mathbf{Hilb}, \quad C \mapsto \mathcal{H}_C. \tag{19}$$

*The restriction maps $S(f) : \mathcal{H}_C \to \mathcal{H}_{C'}$ are isometric (inner-product-preserving) linear maps.*

## 5.3 Uniqueness of the Inner Product

**Proposition 5.5.** *The Hermitian inner product on each fiber $S(C)$ is unique up to an overall positive real scalar (which can be fixed by normalization).*

*Proof.* Suppose $\langle -, - \rangle$ and $\langle -, - \rangle'$ are two Hermitian inner products on $S(C)$ both satisfying the conditions of Theorem 5.3. Then $\langle -, - \rangle' = \langle -, T(-) \rangle$ for a unique positive self-adjoint operator $T$. The restriction compatibility condition forces $T$ to commute with all restriction maps $S(f)$; if the family of restriction maps acts irreducibly on $S(C)$ (which is the case when $\mathscr{C}$ has "enough" morphisms—a condition guaranteed by the Yoneda Constraint), then Schur's lemma implies $T = \lambda \cdot \mathrm{id}$ for some $\lambda > 0$. □

# 6 Observables as Self-Adjoint Natural Transformations

## 6.1 Natural Transformations and Context-Covariance

**Definition 6.1** (Endomorphism of a Presheaf)**.** An **endomorphism** of a physical system $S : \mathscr{C}^{\mathrm{op}} \to \mathbf{Hilb}$ is a natural transformation $\alpha : S \Rightarrow S$. This consists of a family of linear maps $\alpha_C : \mathcal{H}_C \to \mathcal{H}_C$, one for each context $C$, satisfying the naturality condition: for every morphism $f : C' \to C$,

$$\begin{tikzcd} \mathcal{H}_C \arrow[r, "\alpha_C"] \arrow[d, "S(f)"'] & \mathcal{H}_C \arrow[d, "S(f)"] \\ \mathcal{H}_{C'} \arrow[r, "\alpha_{C'}"'] & \mathcal{H}_{C'} \end{tikzcd} \tag{20}$$

That is, $S(f) \circ \alpha_C = \alpha_{C'} \circ S(f)$ for all $f$.

The naturality condition (20) has a transparent physical meaning: an endomorphism $\alpha$ "looks the same from every perspective." If we refine our context from $C$ to $C'$ via $f$, applying $\alpha$ before or after the refinement yields the same result.

## 6.2 Self-Adjointness from Perspectival Consistency

**Definition 6.2** (Observable)**.** An **observable** is an endomorphism $\alpha : S \Rightarrow S$ that is **self-adjoint** with respect to the perspectival inner product: for every context $C$ and all $s, s' \in \mathcal{H}_C$,

$$\langle \alpha_C(s), s' \rangle_C = \langle s, \alpha_C(s') \rangle_C. \tag{21}$$

**Theorem 6.3** (Characterization of Observables)**.** *Let $\alpha : S \Rightarrow S$ be a natural transformation. If $\alpha$ preserves the perspectival pairing in the sense that*

$$\langle \alpha_C(s_1), s_2 \rangle_C = \langle s_1, \alpha_C(s_2) \rangle_C \tag{22}$$

*for all $C, s_1, s_2$, then:*

*(a) $\alpha_C$ is a self-adjoint (Hermitian) operator on $\mathcal{H}_C$ for each $C$.*

*(b) $\alpha_C$ has real eigenvalues.*

*(c) $\mathcal{H}_C$ admits an orthonormal basis of eigenvectors of $\alpha_C$.*

*(d) The eigenvalue decomposition is compatible with restriction maps: if $s$ is an eigenvector of $\alpha_C$ with eigenvalue $\lambda$, then $S(f)(s)$ is an eigenvector of $\alpha_{C'}$ with the same eigenvalue $\lambda$ (or is zero).*

*Proof.* (a) This is immediate from the definition and the fact that $\mathcal{H}_C$ is a complex Hilbert space.
(b) If $\alpha_C(s) = \lambda s$ for $s \neq 0$, then $\lambda \langle s, s \rangle = \langle \alpha_C(s), s \rangle = \langle s, \alpha_C(s) \rangle = \bar{\lambda} \langle s, s \rangle$. Since $\langle s, s \rangle > 0$, we get $\lambda = \bar{\lambda}$, so $\lambda \in \mathbb{R}$.
(c) The spectral theorem for self-adjoint operators on finite-dimensional complex Hilbert spaces gives an orthonormal eigenbasis. (Note: this is where algebraic closure of $\mathbb{C}$ is used—every characteristic polynomial splits over $\mathbb{C}$.)

(d) If $\alpha_C(s) = \lambda s$, then by naturality:

$$\alpha_{C'}(S(f)(s)) = S(f)(\alpha_C(s)) = S(f)(\lambda s) = \lambda \cdot S(f)(s). \tag{23}$$

Hence $S(f)(s)$ is either an eigenvector of $\alpha_{C'}$ with eigenvalue $\lambda$ or is the zero vector. □

## 6.3 The Algebra of Observables

**Proposition 6.4.** *The set of observables on $S$ forms a real Jordan algebra under the symmetrized product*

$$\alpha \circ_J \beta := \frac{1}{2}(\alpha \circ \beta + \beta \circ \alpha) \tag{24}$$

*where $\circ$ denotes composition of natural transformations (applied componentwise at each context).*

*Proof.* The composition of two self-adjoint natural transformations need not be self-adjoint ($(\alpha\beta)^* = \beta^*\alpha^* = \beta\alpha \neq \alpha\beta$ in general). However, the symmetrized product $\alpha \circ_J \beta$ is self-adjoint:

$$(\alpha \circ_J \beta)^* = \frac{1}{2}(\alpha\beta + \beta\alpha)^* = \frac{1}{2}(\beta^*\alpha^* + \alpha^*\beta^*) = \frac{1}{2}(\beta\alpha + \alpha\beta) = \alpha \circ_J \beta. \tag{25}$$

The Jordan identity $(\alpha \circ_J (\beta \circ_J \alpha^2)) = (\alpha \circ_J \beta) \circ_J \alpha^2$ is verified by direct computation.

Naturality is preserved: if $\alpha$ and $\beta$ are natural, then $\alpha \circ_J \beta$ is natural because composition and addition of natural transformations are natural. □

*Remark* 6.5. The Jordan algebra structure of observables was first identified by Jordan, von Neumann, and Wigner [16]. In our framework, it arises as a consequence of the naturality condition combined with self-adjointness, rather than being postulated independently.

## 6.4 Commuting and Non-Commuting Observables

**Definition 6.6** (Compatible Observables)**.** Two observables $\alpha, \beta : S \Rightarrow S$ are **compatible** (or **simultaneously measurable**) if $\alpha_C \circ \beta_C = \beta_C \circ \alpha_C$ for all contexts $C$. They are **complementary** if there is no context $C$ for which $\alpha_C$ and $\beta_C$ share a common eigenbasis.

**Proposition 6.7.** *If $\alpha$ and $\beta$ are complementary observables, then the associated measurement contexts $C_\alpha$ and $C_\beta$ do not admit a common refinement $C_{\alpha\beta}$ such that $S(C_{\alpha\beta})$ simultaneously diagonalizes both $\alpha$ and $\beta$. This is the categorical expression of the Heisenberg uncertainty principle.*

*Proof.* Suppose a common refinement $C_{\alpha\beta}$ existed with restriction maps $f_\alpha : C_{\alpha\beta} \to C_\alpha$ and $f_\beta : C_{\alpha\beta} \to C_\beta$. By naturality, the operators $\alpha_{C_{\alpha\beta}}$ and $\beta_{C_{\alpha\beta}}$ would both be well-defined on $\mathcal{H}_{C_{\alpha\beta}}$, and compatibility with the restriction maps would force them to commute. But complementarity means they share no common eigenbasis, contradicting the spectral theorem applied to two commuting self-adjoint operators. Hence no such common refinement exists. □

# 7 The Born Rule from the Yoneda Isomorphism

This section contains the central derivation: we prove that the probability of a measurement outcome is given by the Born rule $p(\lambda) = |\langle e_\lambda, \psi\rangle|^2$. The proof synthesizes the Yoneda isomorphism with Gleason's theorem, showing that the latter is a structural consequence of the former within the presheaf topos.

## 7.1 The Yoneda Isomorphism in the Physical Setting

Recall the Yoneda Lemma: for any presheaf $F$ on $\mathscr{C}$ and any object $C$,

$$\mathrm{Nat}(\mathsf{y}(C), F) \cong F(C). \tag{26}$$

In our physical setting, with $F = S$ a physical system, this becomes:

**Proposition 7.1** (Physical Yoneda Isomorphism)**.** *The set of natural transformations from the representable presheaf $\mathsf{y}(C) = \mathrm{Hom}_{\mathscr{C}}(-, C)$ to the physical system $S$ is in natural bijection with $S(C)$:*

$$\mathrm{Nat}(\mathsf{y}(C), S) \cong S(C) = \mathcal{H}_C. \tag{27}$$

*The bijection sends a natural transformation $\eta : \mathsf{y}(C) \Rightarrow S$ to the element $\eta_C(\mathrm{id}_C) \in S(C)$.*

*Proof.* This is the Yoneda Lemma. Given $\eta : \mathsf{y}(C) \Rightarrow S$, the element $\eta_C(\mathrm{id}_C) \in S(C)$ determines $\eta$ completely: for any object $C'$ and any $f \in \mathsf{y}(C)(C') = \mathrm{Hom}_{\mathscr{C}}(C', C)$,

$$\eta_{C'}(f) = S(f)(\eta_C(\mathrm{id}_C)) \tag{28}$$

by naturality of $\eta$. Conversely, any element $s \in S(C)$ defines a natural transformation by this formula. □

## 7.2 Physical Interpretation

The physical content of Proposition 7.1 is profound. The natural transformations $\mathrm{Nat}(\mathsf{y}(C), S)$ represent all possible *ways of probing* the system $S$ from the perspective of context $C$. The Yoneda isomorphism says that these probing-ways are in perfect correspondence with the data $S(C)$ that $S$ assigns to $C$. In other words:

> *The ways of looking at a system from a context are exactly the things the system shows to that context.*

This is the perspectival completeness of the Yoneda Constraint: there is no gap between "probing" and "data."

## 7.3 Measurement Contexts and Projection

**Definition 7.2** (Measurement Context)**.** A **measurement context** for an observable $\alpha$ with eigenvalue $\lambda$ is a context $C_{\alpha,\lambda} \in \mathscr{C}$ such that $\mathcal{H}_{C_{\alpha,\lambda}}$ is the eigenspace of $\alpha$ corresponding to $\lambda$. We write $P_\lambda : \mathcal{H} \to \mathcal{H}_{C_{\alpha,\lambda}}$ for the orthogonal projection.

For a measurement context $C_{\alpha,\lambda}$, the Yoneda isomorphism gives

$$\mathrm{Nat}(\mathsf{y}(C_{\alpha,\lambda}), S) \cong S(C_{\alpha,\lambda}) = \mathcal{H}_{C_{\alpha,\lambda}} = \mathrm{im}(P_\lambda). \tag{29}$$

The natural transformation $\eta_\psi : \mathsf{y}(C_{\alpha,\lambda}) \Rightarrow S$ corresponding to a state $\psi$ is determined by $P_\lambda\psi$—the projection of $\psi$ onto the $\lambda$-eigenspace.

## 7.4 The Probability Measure on Outcomes

**Definition 7.3** (Probability Valuation)**.** A **probability valuation** on the outcomes of measuring observable $\alpha$ on system $S$ in state $\psi$ is a function $p : \mathrm{Spec}(\alpha) \to [0, 1]$ (where $\mathrm{Spec}(\alpha)$ is the set of eigenvalues of $\alpha$) satisfying:

(P1) **Non-negativity**: $p(\lambda) \geq 0$ for all $\lambda$.

(P2) **Normalization**: $\sum_{\lambda \in \mathrm{Spec}(\alpha)} p(\lambda) = 1$.

(P3) **Invariance**: $p$ is invariant under the natural automorphisms of $\mathsf{y}(C_{\alpha,\lambda})$ that preserve the inner product—i.e., under the unitary group $U(\mathcal{H}_{C_{\alpha,\lambda}})$.

## 7.5 The Main Theorem

**Theorem 7.4** (Born Rule from Yoneda)**.** *Let $S : \mathscr{C}^{\mathrm{op}} \to$ **Hilb** be a physical system satisfying the Yoneda Constraint, with $\dim(\mathcal{H}_C) \geq 3$ for the relevant contexts. Let $\alpha$ be an observable with eigenvalues $\{\lambda_i\}$ and eigenprojections $\{P_i\}$, and let $\psi \in \mathcal{H}$ be a normalized state. Then the unique probability valuation on the outcomes of $\alpha$ satisfying (P1)–(P3) is*

$$p(\lambda_i) = \langle \psi, P_i \psi \rangle = \|P_i \psi\|^2. \tag{30}$$

*In the non-degenerate case (each eigenspace one-dimensional), this reduces to*

$$p(\lambda_i) = |\langle e_i, \psi \rangle|^2 \tag{31}$$

*where $\{e_i\}$ is an orthonormal eigenbasis of $\alpha$.*

*Proof.* The proof proceeds in four stages.

**Stage 1: From Yoneda to frame functions.** The Yoneda isomorphism $\mathrm{Nat}(\mathsf{y}(C_{\alpha,\lambda_i}), S) \cong \mathcal{H}_{C_{\alpha,\lambda_i}}$ identifies each probe from the measurement context $C_{\alpha,\lambda_i}$ with a vector in the eigenspace. The probability $p(\lambda_i)$ is therefore a function of the relationship between $\psi$ and the eigenspace $\mathcal{H}_{C_{\alpha,\lambda_i}}$—specifically, a function of the projection $P_i \psi$.

The normalization condition (P2) requires $\sum_i p(\lambda_i) = 1$, which means $p$ is determined by its values on the projections $P_i$. Since $\sum_i P_i = \mathrm{id}$ (completeness of eigenprojections), $p$ extends to a function on all rank-one projections (by considering all possible observables), yielding a *frame function*: a function $\mu$ on rank-one projections in $\mathcal{H}$ satisfying $\sum_i \mu(P_{e_i}) = 1$ for every orthonormal basis $\{e_i\}$.

**Stage 2: Invariance from naturality.** The invariance condition (P3) states that $p$ is unchanged under unitary transformations of the eigenspaces. This arises from the naturality of the Yoneda isomorphism: the automorphisms of $\mathsf{y}(C_{\alpha,\lambda_i})$ are precisely the unitaries of $\mathcal{H}_{C_{\alpha,\lambda_i}}$, and naturality ensures that the probability measure transforms covariantly.

More precisely, for any unitary $U \in U(\mathcal{H})$ that permutes the eigenspaces of $\alpha$, the naturality of the Yoneda isomorphism gives $p_U(\lambda_i) = p(U^{-1}\lambda_i)$, where $p_U$ is the probability computed in the rotated basis. The perspectival consistency condition (Theorem 5.3) ensures that the inner product, and hence the probability, is invariant under such unitaries.

**Stage 3: Gleason's theorem.** A frame function $\mu$ on a Hilbert space $\mathcal{H}$ of dimension $n \geq 3$ that is non-negative and sums to 1 on every orthonormal basis is necessarily of the form $\mu(P_e) = \mathrm{Tr}(\rho P_e)$ for a unique positive trace-class operator $\rho$ with $\mathrm{Tr}(\rho) = 1$ (a density operator). This is Gleason's theorem [14].

In our setting, Gleason's theorem is not an independent input but is *forced* by the structure of the presheaf topos. The lattice of closed subspaces of $\mathcal{H}_C$ is isomorphic to the lattice of subobjects of $S|_C$ in the presheaf topos restricted to measurement contexts. The Yoneda Constraint requires that probabilities be defined on this lattice in a way compatible with the natural transformations, and the non-Boolean structure of the subobject lattice (a consequence of the non-commutativity of complementary observables—Proposition 6.7) eliminates all dispersion-free probability assignments.

The dimension condition $\dim(\mathcal{H}) \geq 3$ is essential: in dimension 2, there exist "exotic" frame functions not of the Gleason form. The physical interpretation is that a system with at least three distinguishable outcomes for some observable has a rich enough perspectival structure to force the Born rule.

**Stage 4: Pure-state reduction.** For a pure state $\psi$ (i.e., a state that is an extremal point of the convex set of density operators), $\rho = |\psi\rangle\langle\psi|$, and Gleason's formula reduces to

$$p(\lambda_i) = \mathrm{Tr}(|\psi\rangle\langle\psi|P_i) = \langle\psi, P_i\psi\rangle = \|P_i\psi\|^2. \tag{32}$$

For non-degenerate eigenvalues, $P_i = |e_i\rangle\langle e_i|$, giving $p(\lambda_i) = |\langle e_i, \psi\rangle|^2$. □

*Remark* 7.5. The derivation reveals the precise role of each ingredient:

- The **Yoneda isomorphism** connects probes to data, establishing that probabilities are functions of projections.
- **Perspectival consistency** (the inner product) provides the quadratic structure.
- **Gleason's theorem** eliminates all alternatives to the $|\langle e, \psi\rangle|^2$ form, given the lattice structure.
- The **naturality** condition ensures unitary invariance, which is one of Gleason's hypotheses.

The Born rule is therefore not a separate postulate but a structural theorem of the presheaf topos.

## 7.6 The Two-Dimensional Case

**Proposition 7.6** (Monoidal Bootstrap for Qubits)**.** *For* $\dim(\mathcal{H}) = 2$ *(qubit systems), Gleason's theorem does not apply directly. However, the monoidal structure of* $\mathscr{C}$ *provides a bootstrap:*

*(a) The composite of two qubits lives in* $\mathcal{H} \otimes \mathcal{H}$*, which has dimension* $4 \geq 3$*.*

*(b) Gleason's theorem applies to* $\mathcal{H} \otimes \mathcal{H}$*, establishing the Born rule for the composite system.*

*(c) The Born rule for the individual qubit follows by restriction (partial trace):* $p_A(\lambda) = \mathrm{Tr}_B(\rho_{AB} \cdot (P_\lambda \otimes \mathrm{id}_B))$*.*

*Thus the Yoneda Constraint, through the monoidal structure, bootstraps the Born rule even for two-dimensional systems.*

*Proof.* The monoidal presheaf condition gives $S(C_A \otimes C_B) \cong S(C_A) \otimes S(C_B) = \mathcal{H}_A \otimes \mathcal{H}_B$, with $\dim(\mathcal{H}_A \otimes \mathcal{H}_B) = 4$. Theorem 7.4 applies to the composite, giving $p(\lambda_i, \mu_j) = \langle \Psi, (P_i^A \otimes P_j^B)\Psi \rangle$ for joint outcomes. The marginal probability for subsystem $A$ is $p_A(\lambda_i) = \sum_j p(\lambda_i, \mu_j) = \langle \Psi, (P_i^A \otimes \mathrm{id}_B)\Psi \rangle = \mathrm{Tr}_B(\rho_{AB} \cdot (P_i^A \otimes \mathrm{id}_B))$. For a product state $\Psi = \psi_A \otimes \psi_B$, this reduces to $p_A(\lambda_i) = |\langle e_i, \psi_A \rangle|^2$—the Born rule for the qubit. □

# 8 Superposition as Perspectival Richness

## 8.1 The Standard Puzzle of Superposition

In the standard formulation, a superposition $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ is a state that is "simultaneously" in two different eigenstates of some observable. This description leads to the notorious measurement problem: if the state is simultaneously $|0\rangle$ and $|1\rangle$, why do we observe only one outcome?

The Yoneda framework dissolves this puzzle by reinterpreting superposition as *perspectival richness*: the system has definite data in each context, but different contexts yield different data, and no single context captures the full content of the presheaf.

## 8.2 Formal Characterization

**Theorem 8.1** (Superposition as Perspectival Indefiniteness)**.** *Let $\alpha$ be an observable with eigenstates $\{|e_i\rangle\}$ and let $|\psi\rangle = \sum_i c_i|e_i\rangle$ with at least two nonzero coefficients. Then:*

*(a) In the measurement context $C_\alpha$, the presheaf assigns $\psi$ as a vector in $\mathcal{H}_{C_\alpha}$ that is not an eigenvector of $\alpha_{C_\alpha}$.*

*(b) For each eigenvalue $\lambda_j$, the restriction to the measurement context $C_{\alpha,\lambda_j}$ yields $P_j|\psi\rangle = c_j|e_j\rangle$—a definite element of $\mathcal{H}_{C_{\alpha,\lambda_j}}$.*

*(c) For a complementary observable $\beta$, the state $|\psi\rangle$ may be an eigenvector of $\beta$, yielding definite data in context $C_\beta$.*

*(d) The totality of these context-dependent data—some definite, some indefinite—constitutes the presheaf $S_\psi$, and the Yoneda Lemma guarantees that $|\psi\rangle$ is completely determined by this totality.*

*Proof.* (a) By definition, $\alpha_{C_\alpha}|\psi\rangle = \sum_i c_i\lambda_i|e_i\rangle \neq \lambda_j|\psi\rangle$ for any single $j$ (since at least two $c_i$ are nonzero and the eigenvalues are distinct). Hence $|\psi\rangle$ is not an eigenvector.
(b) The projection $P_j$ picks out the $\lambda_j$-component: $P_j|\psi\rangle = c_j|e_j\rangle$, which is a well-defined (possibly zero) element of the eigenspace. The restriction map $S(f_j) : \mathcal{H}_{C_\alpha} \to \mathcal{H}_{C_{\alpha,\lambda_j}}$, where $f_j : C_{\alpha,\lambda_j} \to C_\alpha$ is the inclusion of the eigenspace context, is exactly $P_j$.
(c) If $|\psi\rangle$ happens to be an eigenvector of $\beta$, then $\beta_{C_\beta}|\psi\rangle = \mu|\psi\rangle$ for some eigenvalue $\mu$, and $|\psi\rangle$ presents definite data in $C_\beta$.
(d) The presheaf $S_\psi$ assigns to each context $C$ the vector $S_\psi(C) =$ "the appearance of $|\psi\rangle$ in context $C$," computed by the appropriate restriction maps. By the Yoneda Lemma, $|\psi\rangle \cong S_\psi$ as objects of $\widehat{\mathscr{C}}$ (more precisely, the natural transformation corresponding to $|\psi\rangle$ via the Yoneda isomorphism encodes all this data). Full faithfulness of $\mathsf{y}$ ensures no information is lost. □

## 8.3 The Mountain Analogy, Made Precise

The presheaf picture of superposition is analogous to a mountain that looks different from different vantage points. But the analogy is now mathematically precise:

- The **mountain** is the presheaf $S_\psi$.
- The **vantage points** are the objects $C \in \mathscr{C}$.
- The **appearance from each vantage point** is $S_\psi(C) \in \mathcal{H}_C$.
- The **Yoneda Lemma** guarantees that the mountain is completely determined by the totality of appearances.
- A **superposition** is a mountain that looks "sharp" from some vantage points (eigenbasis contexts) and "blurry" from others (non-eigenbasis contexts).

There is nothing paradoxical about this. The "paradox" of superposition arises only if one insists that the system must have a definite state *independent of context*—i.e., if one denies the Yoneda Constraint.

## 8.4 Non-Existence of Dispersion-Free States

**Corollary 8.2.** *There is no state $|\psi\rangle$ that is simultaneously an eigenvector of all observables—i.e., no state presents definite data to every context. This is the Kochen–Specker theorem, derived here as a consequence of the non-Boolean structure of the subobject lattice in $\widehat{\mathscr{C}}$.*

*Proof.* Suppose $|\psi\rangle$ were simultaneously an eigenvector of all observables. Then the presheaf $S_\psi$ would assign a *global section* to the subobject lattice—a consistent assignment of eigenvalues to all observables simultaneously. But the subobject lattice in the presheaf topos is non-Boolean (it does not satisfy the distributive law), and non-Boolean lattices of dimension $\geq 3$ do not admit dispersion-free probability measures (Kochen–Specker [15]; see also Bell [20] and the original quantum logic of Birkhoff and von Neumann [21]). Hence no such $|\psi\rangle$ exists. $\square$

# 9 Comparison with Other Derivations

We now compare our derivation systematically with the three major programs that derive quantum theory from operational axioms.

## 9.1 Hardy's Five Axioms (2001)

Hardy [8] derives quantum theory from five "reasonable" axioms:

1. **Probabilities**: Relative frequencies converge.
2. **Simplicity**: The number of parameters $K$ needed to describe a state is minimal for a given number of distinguishable outcomes $N$. Specifically, $K$ is a function of $N$ only.
3. **Subspaces**: Systems with $M \leq N$ distinguishable outcomes behave as subsystems.
4. **Composite systems**: Joint states of independent systems are determined by local measurements and correlations.
5. **Continuity**: There exists a continuous reversible transformation between any two pure states.

**Proposition 9.1** (Subsumption of Hardy's Axioms)**.** *Each of Hardy's axioms is either a consequence of the Yoneda Constraint applied to the category $\mathscr{C}$, or an additional specification of $\mathscr{C}$:*

- ***Axiom 1****: Probabilities emerge from the Born rule (Theorem 7.4), which is derived rather than assumed.*
- ***Axiom 2****: The "simplicity" condition $K = N^2 - 1$ for complex quantum theory follows from the dimension of the space of density matrices on $\mathbb{C}^N$, which is forced by our linearization (Theorem 3.4) and the choice $k = \mathbb{C}$ (Theorem 4.4).*
- ***Axiom 3****: The subspace axiom corresponds to the existence of sub-presheaves and the compatibility of the linear structure with restriction maps.*
- ***Axiom 4****: The composite system axiom is our monoidal presheaf condition (Definition 3.1).*
- ***Axiom 5****: Continuity of transformations is encoded in the topological structure of $\mathrm{Hom}_{\mathscr{C}}(C, C')$ and the naturality of automorphisms.*

The key advantage of the Yoneda approach over Hardy's is that Hardy's "simplicity" axiom—arguably his most controversial—is derived rather than assumed. It follows from the structural properties of presheaves on a monoidal category, without needing to postulate a specific relationship between $K$ and $N$.

## 9.2 Chiribella–D'Ariano–Perinotti (2011)

Chiribella, D'Ariano, and Perinotti [9] derive quantum theory from six axioms within the framework of operational-probabilistic theories (OPTs):

1. **Causality**: The probability of a preparation is independent of the choice of future test.
2. **Perfect distinguishability**: Every state has at least one state from which it is perfectly distinguishable.
3. **Ideal compression**: Information can be encoded losslessly in a system of minimal dimension.
4. **Local distinguishability**: Global states are determined by local measurements and correlations.
5. **Pure conditioning**: After a measurement with a pure outcome, the remaining system is in a pure state.
6. **Purification**: Every mixed state of $A$ arises as the marginal of a pure state of $AB$, and this purification is essentially unique.

**Proposition 9.2** (Subsumption of CDP Axioms)**.** *The CDP axioms relate to the Yoneda framework as follows:*

- ***Causality**: In the presheaf framework, causality is the statement that the restriction map $S(f) : S(C_{future}) \to S(C_{past})$ depends only on $f$ and not on other morphisms from $C_{past}$. This is automatic from presheaf functoriality.*
- ***Local distinguishability**: This is exactly the monoidal presheaf condition $S(C \otimes C') \cong S(C) \otimes S(C')$ (over* **Hilb***, this becomes local tomography).*
- ***Purification**: In the presheaf topos, every object has an injective hull (an analogue of purification), and the uniqueness of purification corresponds to the essential uniqueness of the injective resolution in the topos. This is a deep structural property of presheaf topoi rather than an independent axiom.*
- ***Pure conditioning, perfect distinguishability, ideal compression**: These follow from the Hilbert-space structure of the fibers and the properties of orthogonal projections established in Theorems 5.3 and 6.3.*

The main advance over CDP is the replacement of six operational axioms—each requiring separate justification—with a single structural principle. The "purification" axiom, which CDP identify as the most characteristically quantum postulate, emerges as a structural theorem of the presheaf topos.

## 9.3 Masanes–Müller (2011)

Masanes and Müller [10] derive finite-dimensional quantum theory from five postulates:

1. In systems that carry one bit of information, each state is characterized by a finite set of outcome probabilities.
2. The state of a composite system is characterized by the statistics of local measurements.
3. All systems that effectively carry the same amount of information have equivalent state spaces.
4. Any pure state of a system can be reversibly transformed into any other.
5. In systems that carry one bit of information, all mathematically well-defined measurements are allowed.

**Proposition 9.3.** *The Masanes–Müller postulates are related to our framework as follows:*

- ***Postulate 1****: Follows from the finite-dimensionality of $\mathcal{H}_C$ (the finiteness condition in Theorem 3.4).*
- ***Postulate 2****: Is the monoidal presheaf condition.*
- ***Postulate 3****: Corresponds to the functoriality of $S$—isomorphic contexts yield isomorphic fibers.*
- ***Postulate 4****: Is the transitivity of the unitary group action on pure states, which follows from the Hilbert-space structure.*
- ***Postulate 5****: Is the statement that every self-adjoint natural transformation is a legitimate observable (Theorem 6.3).*

## 9.4 Summary of Comparison

| **Feature** | **Yoneda** | **Hardy** | **CDP** | **MM** |
|---|---|---|---|---|
| Number of axioms | 1 | 5 | 6 | 5 |
| Hilbert space derived | Yes | Yes | Yes | Yes |
| $\mathbb{C}$ derived | Yes | Yes* | No† | Yes |
| Born rule derived | Yes | Partial | Yes | Partial |
| Superposition explained | Yes | No | No | No |
| Measurement problem addressed | Yes | No | No | No |
| Category-theoretic | Yes | No | No | No |

* Via the simplicity axiom. † CDP assume $\mathbb{C}$ from the start.

The Yoneda approach is distinguished by its economy (a single axiom), its explanatory power (superposition and measurement have transparent interpretations), and its mathematical naturality (it leverages the deepest result of category theory rather than *ad hoc* information-theoretic postulates).

# 10 The Derivation Chain: Summary and Logical Structure

We collect the complete logical flow of the derivation for clarity.

## 10.1 Axiom

We posit a single axiom:

> **Yoneda Constraint.** A physical system is completely characterized by its relational profile across all observational contexts.

## 10.2 Structural Hypotheses on $\mathscr{C}$

We assume that the category $\mathscr{C}$ of observational contexts has:

(H1) A symmetric monoidal product $\otimes$ (parallel combination).

(H2) Finite coproducts $\sqcup$ (exclusive choice).

(H3) Distributivity of $\otimes$ over $\sqcup$.

(H4) A braided monoidal structure with a super-braiding (accommodating fermions).

(H5) Sufficient richness (enough objects and morphisms to make the Yoneda Constraint non-trivial).

## 10.3 The Derivation

| Step | Input | $\Longrightarrow$ | Output |
|---|---|---|---|
| 1 | Yoneda Constraint | $\Longrightarrow$ | Systems are presheaves $S : \mathscr{C}^{\mathrm{op}} \to \mathbf{Set}$ |
| 2 | $\otimes + \sqcup +$ distributivity | $\Longrightarrow$ | $S(C)$ is a vector space (Thm. 3.4) |
| 3 | Cancellation (interference) | $\Longrightarrow$ | $S(C)$ is a module over a division ring |
| 4 | Super-braiding | $\Longrightarrow$ | Division ring is $\mathbb{C}$ (Thm. 4.4) |
| 5 | Perspectival consistency | $\Longrightarrow$ | Hermitian inner product (Thm. 5.3) |
| 6 | Naturality + self-adjointness | $\Longrightarrow$ | Observables (Thm. 6.3) |
| 7 | Yoneda iso. + Gleason | $\Longrightarrow$ | Born rule (Thm. 7.4) |
| 8 | Presheaf structure | $\Longrightarrow$ | Superposition (Thm. 8.1) |

## 10.4 What Is and Is Not Derived

**Derived from the Yoneda Constraint + structural hypotheses on $\mathscr{C}$:**

- Complex Hilbert space structure of state spaces.
- Self-adjoint operators as observables.
- The Born rule $p(\lambda) = |\langle e_\lambda, \psi\rangle|^2$.

- The superposition principle.
- The uncertainty principle for complementary observables.
- The Kochen–Specker theorem (non-existence of dispersion-free states).

**Not derived (requires additional input):**

- The specific Hilbert space $\mathcal{H}$ for a given physical system (this depends on the choice of $\mathscr{C}$ and the presheaf $S$).
- The Hamiltonian (dynamics requires specifying a one-parameter family of natural automorphisms).
- The particle content and gauge symmetries of the Standard Model.
- Quantum field theory and the passage to infinite dimensions.

The structural hypotheses (H1)–(H5) are themselves physically motivated—they describe features of any reasonable collection of experimental contexts—but they are not derived from the Yoneda Constraint alone. They should be viewed as structural properties of $\mathscr{C}$ rather than independent physical axioms.

# 11 Discussion and Open Problems

## 11.1 The Status of the Derivation

We have shown that the Yoneda Constraint, applied to a category of observational contexts with monoidal and coproduct structure, forces the emergence of the core mathematical structures of quantum mechanics. The derivation is not a reformulation or restatement of quantum mechanics in categorical language; it is a *derivation*—the quantum formalism emerges as a theorem from premises that make no reference to Hilbert spaces, operators, or probabilities.

The single most important conceptual achievement is the treatment of the Born rule. In every other axiomatization, the Born rule (or an equivalent probabilistic postulate) must be assumed or motivated by hand. Here, it emerges from the interplay of three purely structural features: the Yoneda isomorphism, the inner product from perspectival consistency, and the non-Boolean structure of the presheaf lattice (which triggers Gleason's theorem).

## 11.2 Relation to Topos Quantum Theory

Our approach has deep connections to the topos-theoretic quantum mechanics of Isham, Butterfield, Döring, and others [4, 5, 6]. The key difference is in the direction of reasoning. The topos program begins with the standard quantum formalism and seeks to reformulate it within a presheaf topos (typically over the category of commutative subalgebras of a von Neumann algebra). We proceed in the opposite direction: we begin with the presheaf structure and *derive* the quantum formalism.

Our category $\mathscr{C}$ of observational contexts is related to, but distinct from, the category of commutative subalgebras used in the Döring–Isham program. In our framework, the objects of $\mathscr{C}$ are physically motivated (experimental contexts), while in the Döring–Isham framework, the objects are algebraically motivated (commutative subalgebras). A natural conjecture is that the two categories are equivalent for suitable choices of $\mathscr{C}$.

## 11.3 The Measurement Problem Revisited

In the Yoneda framework, the measurement problem does not arise in its traditional form. There is no ontological collapse: the presheaf assigns definite data to each context, and "measurement" is simply the selection of a particular context. The appearance of a single definite outcome in context $C$ is the restriction $S(f) : S(C') \to S(C)$ for the morphism $f$ corresponding to the experimental intervention.

This resolution is closest in spirit to relational quantum mechanics (RQM) [7], but with a crucial difference: in RQM, the relational nature of quantum states is posited philosophically; in our framework, it is a mathematical theorem (the Yoneda Lemma). The broader program of Quantum Perspectivism, including its implications for quantum gravity and emergent spacetime, is developed in [25, 26].

## 11.4 Open Problems

1. **Infinite-dimensional extension.** Our derivation assumes finite-dimensional Hilbert spaces. Extending to the infinite-dimensional case (relevant for quantum

field theory) requires working with $\mathbf{Hilb}_\infty$-valued presheaves and replacing Gleason's theorem with its infinite-dimensional generalization.

2. **Determination of $\mathscr{C}$ from first principles.** We have left the category $\mathscr{C}$ partially unspecified. Determining $\mathscr{C}$ from physical first principles—possibly from the causal structure of spacetime—is a major open problem.

3. **Dynamics.** We have not derived the Schrödinger equation. The one-parameter family of natural automorphisms generating time evolution needs to be connected to the structure of $\mathscr{C}$ (possibly through a temporal ordering on contexts).

4. **Quantum gravity.** The Yoneda framework naturally suggests that spacetime itself is a derived concept, emerging from the Grothendieck topology on $\mathscr{C}$. Developing this into a complete quantum gravity program is the most ambitious open direction.

5. **Experimental tests.** Can the Yoneda framework make predictions that differ from standard quantum mechanics? Possible avenues include exotic probability rules in the two-dimensional case (before the monoidal bootstrap), modifications at the Planck scale, and novel entanglement structures.

# 12 Conclusion

We have derived the core mathematical structures of quantum mechanics—complex Hilbert spaces, the inner product, self-adjoint observables, and the Born rule—from a single foundational principle: the Yoneda Constraint. The derivation requires structural hypotheses on the category of observational contexts (monoidal and coproduct structure, braiding, distributivity) that are physically natural and categorically mild.

The key results are:

(i) **Linearization** (Theorem 3.4): The interplay of parallel combination and exclusive choice of contexts forces vector-space structure on the fibers of any monoidal presheaf.

(ii) **Complex numbers** (Theorem 4.4): The super-braiding that accommodates fermionic statistics requires the ground field to contain $\sqrt{-1}$, selecting $\mathbb{C}$ as the unique commutative, algebraically closed field of characteristic zero.

(iii) **Inner product** (Theorem 5.3): Perspectival consistency—the requirement that data from different contexts be coherently related through common refinements—uniquely determines a Hermitian inner product on each fiber.

(iv) **Observables** (Theorem 6.3): Context-covariant endomorphisms of the presheaf that respect the inner product are precisely self-adjoint operators, with real eigenvalues and orthogonal eigenbases.

(v) **Born rule** (Theorem 7.4): The Yoneda isomorphism, combined with the inner product and Gleason's theorem, uniquely determines the probability of measurement outcomes as $p(\lambda) = |\langle e_\lambda, \psi\rangle|^2$.

(vi) **Superposition** (Theorem 8.1): A state in superposition is a presheaf that presents different definite data to different contexts, with no single context capturing the complete information.

The philosophical upshot is that quantum mechanics is not an empirical discovery that might have been otherwise. It is the unique physics consistent with the Yoneda Lemma—the mathematical theorem that identity is relational structure. If physical systems have no intrinsic properties beyond their relationships to observational contexts, then the world must be quantum.

---

**Acknowledgments.** The author thanks the YonedaAI Research Collective for ongoing collaboration and intellectual support. This work builds on the foundational contributions of Saunders Mac Lane, Alexander Grothendieck, and Nobuo Yoneda.

**GrokRxiv DOI:** `10.48550/GrokRxiv.2026.02.yoneda-constraint-quantum-structure`

# References

[1] N. Yoneda, "On the homology theory of modules," *J. Fac. Sci. Univ. Tokyo Sect. I*, vol. 7, pp. 193–227, 1954.

[2] S. Mac Lane, *Categories for the Working Mathematician*, 2nd ed., Springer, 1998.

[3] A. Grothendieck, "Sur quelques points d'algèbre homologique," *Tôhoku Math. J.*, vol. 9, pp. 119–221, 1957.

[4] C. J. Isham, "Topos theory and consistent histories: The internal logic of the set of all consistent sets," *Int. J. Theor. Phys.*, vol. 36, pp. 785–814, 1997.

[5] J. Butterfield and C. J. Isham, "A topos perspective on the Kochen–Specker theorem: I. Quantum states as generalized valuations," *Int. J. Theor. Phys.*, vol. 37, pp. 2669–2733, 1998.

[6] A. Döring and C. J. Isham, "A topos foundation for theories of physics," *J. Math. Phys.*, vol. 49, 053515, 2008.

[7] C. Rovelli, "Relational quantum mechanics," *Int. J. Theor. Phys.*, vol. 35, pp. 1637–1678, 1996.

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

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

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

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

[12] H. Barnum and A. Wilce, "Information processing in convex operational theories," *Electron. Notes Theor. Comput. Sci.*, vol. 270, pp. 3–15, 2011.

[13] S. Abramsky and B. Coecke, "A categorical semantics of quantum protocols," in *Proc. 19th Annual IEEE Symposium on Logic in Computer Science*, pp. 415–425, 2004.

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

[15] S. Kochen and E. P. Specker, "The problem of hidden variables in quantum mechanics," *J. Math. Mech.*, vol. 17, pp. 59–87, 1967.

[16] P. Jordan, J. von Neumann, and E. Wigner, "On an algebraic generalization of the quantum mechanical formalism," *Ann. of Math.*, vol. 35, pp. 29–64, 1934.

[17] M. P. Solèr, "Characterization of Hilbert spaces by orthomodular spaces," *Comm. Algebra*, vol. 23, pp. 219–243, 1995.

[18] E. C. G. Stueckelberg, "Quantum theory in real Hilbert space," *Helv. Phys. Acta*, vol. 33, pp. 727–752, 1960.

[19] J. M. Leinaas and J. Myrheim, "On the theory of identical particles," *Nuovo Cimento B*, vol. 37, pp. 1–23, 1977.

[20] J. S. Bell, "On the Einstein Podolsky Rosen paradox," *Physics Physique Fizika*, vol. 1, pp. 195–200, 1964.

[21] G. Birkhoff and J. von Neumann, "The logic of quantum mechanics," *Ann. of Math.*, vol. 37, pp. 823–843, 1936.

[22] E. Riehl, *Category Theory in Context*, Dover Publications, 2016.

[23] P. T. Johnstone, *Sketches of an Elephant: A Topos Theory Compendium*, Oxford University Press, 2002.

[24] J. Ladyman and D. Ross, *Every Thing Must Go: Metaphysics Naturalized*, Oxford University Press, 2007.

[25] M. Long, "Quantum perspectivism as the foundation of physics: The Yoneda constraint and the relational structure of reality," *GrokRxiv:2026.02.quantum-perspectivism-yoneda*, 2026.

[26] M. Long, "Emergent spacetime from quantum information structures," *GrokRxiv*, 2026.