Isosymmetry: Why Quantum Structure Is Independent of Physical Realisation

Keith Taylor
VERSF Theoretical Physics Program

Abstract

Quantum theory exhibits exact structural universality: photons, atoms, spins, and superconducting circuits all obey identical kinematic rules, compositional laws, and correlation constraints. Standard formulations treat this universality as axiomatic. Symmetry principles explain invariance within a given state space but not why the same state space applies across physical domains. We introduce *isosymmetry*: an equivalence principle defined over admissible constraint structures. Two systems are isosymmetric if they exhibit identical patterns of state discrimination, composition, and irreversible outcome production, regardless of microphysical implementation. We show that Hilbert space structure, tensor-product composition, and entanglement universality arise as invariants of isosymmetric constraint classes. Isosymmetry thus explains substrate-independence at a level prior to dynamics, complementing reconstruction programs that derive quantum formalism from operational axioms.

General Reader Abstract

Why do completely different physical systems—light waves, spinning electrons, atoms, and electrical circuits—all follow exactly the same quantum rules? This paper addresses that puzzle.

Usually, physics textbooks simply declare that quantum mechanics works the same way for everything. But that's not an explanation—it's just restating the mystery. We propose a deeper answer: these systems are *isosymmetric*, meaning they can perform the same fundamental tasks with the same resource costs.

Think of it like this: a calculator and a smartphone are made of completely different components, but if they can do the same arithmetic operations with similar effort, they belong to the same "computational class." Similarly, photons and electrons belong to the same "isosymmetric class" because they support the same patterns of distinguishing states, combining systems, and producing definite measurement outcomes.

The key insight is that quantum structure—the mathematical framework of Hilbert spaces, superposition, and entanglement—isn't an arbitrary choice. It's the *only* framework that can consistently describe systems with these operational capabilities. We show this by connecting to rigorous "reconstruction theorems" that derive quantum mechanics from basic operational principles.

This perspective also explains why classical physics is different: classical systems allow unlimited precision in distinguishing states, violating the finite-capacity requirement that defines quantum isosymmetry classes. The quantum-classical divide isn't mysterious—it reflects a fundamental structural difference in what tasks each type of system can perform.

Table of Contents

---

1. Introduction: The Universality Problem

Quantum theory exhibits a striking and unexplained form of universality. Physical systems with radically different microphysical substrates—photons, atomic orbitals, nuclear spins, superconducting circuits, and vibrational modes—are all described by the same abstract structure: complex Hilbert spaces, tensor-product composition, entanglement, and probabilistic measurement outcomes governed by identical rules. This universality is exact, not approximate. The compositional structure, probability calculus, and correlation behaviour do not vary with the underlying physical realisation.

In standard formulations of quantum mechanics, this universality is treated as a kinematic postulate. One assumes a Hilbert space representation for all systems, introduces observables as operators, and proceeds directly to dynamics and measurement. While enormously successful operationally, this approach leaves unanswered a foundational question: *why does quantum theory not depend on what a system is made of?*

Symmetry principles are often invoked as explanations for universality in physics. Continuous symmetries underpin conservation laws, gauge symmetries organise interactions, and spacetime symmetries constrain dynamics. However, these principles operate *within* a given state space; they do not explain why the same state-space structure applies *across* distinct physical domains.

Recent work in quantum foundations has increasingly emphasised relational and entanglement-first perspectives, suggesting that spacetime locality and even kinematic structure may be emergent rather than fundamental [1–4]. Reconstruction programs have derived quantum formalism from operational axioms [5–8]. These developments sharpen, rather than resolve, the universality problem: they show *what* the structure must be, but not *why* disparate physical systems all satisfy the same axioms.

In this paper, we introduce *isosymmetry*: an equivalence principle defined over admissible constraint structures. Two systems are isosymmetric if they admit the same global constraint resolution behaviour under irreversible transitions, regardless of their microphysical implementation. We argue that Hilbert space structure, tensor composition, and entanglement universality arise as invariants of isosymmetric constraint classes.

We do not assume all fact-producing systems are isosymmetric. We claim: whenever two systems realise the same admissible task class (Definition 4.0), they must share the same kinematic and compositional structure. The observed universality of quantum mechanics reflects the empirical fact that many disparate physical substrates realise the same constraint class.

Explanatory advance. Reconstruction programmes show that if physical systems satisfy certain operational axioms, then the quantum formalism follows. Isosymmetry explains *why* disparate substrates satisfy the same axiom family: because they realise the same admissible task class under finite distinguishability and irreversible record formation. The universality of Hilbert space is thereby relocated from a postulate about "what states are" to a structural fact about which global constraint-resolution tasks are physically admissible.

---

2. The Limits of Symmetry-Based Explanations

Symmetry occupies a central role in modern theoretical physics. Noether's theorem links continuous symmetries to conserved quantities, gauge symmetries organise the interactions of the Standard Model, and spacetime symmetries constrain relativistic dynamics.

However, symmetry-based explanations operate on a pre-existing state space. A symmetry specifies transformations that leave certain relations invariant, but it does not determine the structure of the space on which those transformations act.

This limitation becomes apparent when addressing quantum universality. The question is not why certain quantities are conserved, but why all quantum systems share the same kinematic and compositional framework. Gauge symmetries presuppose quantum structure rather than explain it. Spacetime symmetries cannot ground universality, since quantum theory applies equally in nonrelativistic and relativistic regimes.

Renormalisation and effective field theory explain robustness under perturbation and scale transformation, not identity across realisations. To explain why quantum mechanics is substrate-blind, a different principle is required—one that operates at the level of *what systems can do* rather than *how they transform*.

---

3. Admissibility and Constraint Classes

Physical theories distinguish implicitly between what is mathematically conceivable and what is physically admissible. While mathematical formalisms permit a vast space of possible states and evolutions, only a subset can be realised within a physical universe subject to finite resources, irreversibility, and consistency constraints. This distinction is usually left implicit in quantum mechanics, but it plays a central role in explaining universality.

3.1 Operational Definition of Admissibility

Definition 3.1 (Operational Admissibility). Fix a bounded laboratory domain D and a class of implementable instruments Instr(D), consisting of preparations, transformations, and readout channels that respect:

- (i) finite distinguishability: at most *n* states can be perfectly discriminated in a single measurement,
- (ii) irreversible commitment: measurement outcomes are recorded in stable, non-retractable form, and
- (iii) resource boundedness: operations consume finite resources, bounded by a resource vector $R = (E, \tau, M, \Delta S)$ denoting energy, time, memory, and entropy export. Admissibility is defined relative to bounded regions of this resource space; all admissible tasks are required to have finite R for fixed accuracy targets.

A state or process is *admissible* if and only if it can be realised by some element of Instr(D) without violating these constraints.

This definition grounds admissibility in operational capabilities rather than in abstract mathematical structure. A state is physically admissible if it can be prepared, evolved, and measured using finite physical resources with irreversible outcome registration.

3.2 Constraint Classes

From this perspective, a physical system is characterised not solely by its degrees of freedom or Hamiltonian, but by the class of constraints it admits and the manner in which those constraints can be resolved. These constraints include limits on distinguishability, requirements of global consistency, and the existence of critical thresholds beyond which reversible evolution gives way to irreversible commitment.

Definition 3.2 (Constraint Class). A *constraint class* $\mathcal{C}$ is the set of admissible global states and transitions that a system can support under the conditions of Definition 3.1. Constraint classes abstract away from microphysical details and retain only the structural features relevant to admissibility and resolution.

Two systems may differ radically in their physical composition yet belong to the same constraint class if they support identical patterns of admissible correlation, composition, and resolution.

3.3 Global Character of Admissibility

Crucially, constraint classes are defined independently of spacetime localisation. In entangled systems, admissibility applies to the global configuration rather than to local subsystems. The relevant constraints govern the joint state, not separable components, and resolution acts on the system as a whole. This global character of admissibility is essential for understanding both entanglement universality and the absence of signal-based explanations for correlated outcomes.

3.4 Reversible and Irreversible Regimes

Admissibility introduces a natural distinction between reversible and irreversible regimes. Below certain thresholds, systems may explore admissible states reversibly, maintaining underdetermined or oscillatory configurations. At critical thresholds, however, admissibility constraints force resolution: a previously open set of alternatives collapses into a single committed outcome. The existence of such thresholds is not a dynamical assumption but a structural requirement for fact production.

---

4. Isosymmetry: Formal Statement and Definition

We now introduce the central concept of this paper. The universality of quantum structure across disparate physical systems can be explained if such systems are equivalent at the level that matters for admissibility and resolution, even when they differ microscopically.

4.1 Isosymmetry Distinguished from Symmetry

Isosymmetry is not a symmetry in the conventional sense. It does not refer to a group of transformations acting on states, nor to invariance under continuous or discrete operations. Instead, it is an equivalence relation defined over constraint classes. Where symmetry concerns invariance of form under transformation, isosymmetry concerns invariance of admissible resolution behaviour across realisations.

4.2 Formal Definition

We define isosymmetry first at the level of tasks and instruments, without presupposing a state-space representation.

Definition 4.0 (Task Isosymmetry). Two systems A and B are *isosymmetric* if for every admissible experiment in A—i.e., any finite composition of preparations, transformations, and readouts drawn from Instr_A(D) with resource budget R—there exists a corresponding experiment in B drawn from Instr_B(D) with comparable budget $\|R'\| \leq c\|R\|$ for some fixed constant c and suitable norm on the resource vector, such that the full outcome statistics agree up to relabelling, and vice versa.

This definition grounds isosymmetry in operational procedures rather than abstract mathematical structures. Two systems are isosymmetric if they can perform the same tasks with the same statistical outcomes.

Definition 4.1 (State-Space Isosymmetry; representation form). If the operational statistics of Definition 4.0 admit convex state representations, then isosymmetry induces an operational isomorphism $\varphi: \mathcal{S}\_A \rightarrow \mathcal{S}\_B$ between their admissible state spaces (i.e., an affine bijection preserving convex mixtures and implementable task structure)—where $\mathcal{S}$ denotes the operational (convex) state space of admissible preparations induced by Instr(D)—and the following invariants agree:

- (i) **Discrimination capacity:** max $|\{s \in \mathcal{S} : s \text{ perfectly distinguishable}\}|$ is equal for A and B,

- (ii) **Non-signalling correlation bounds:** for every finite Bell-type scenario (m_A, m_B, k_A, k_B), the set of achievable conditional distributions $P(a,b|x,y)$ under admissible instruments with the same resource bounds (Definition 3.1) is identical up to relabelling, and
- (iii) **Record capacity scaling:** irreversible commitment requires the same resource scaling (bits, time, entropy export). Specifically, record capacity scaling is measured by the minimal entropy export (or equivalent memory resource) required to stabilise a classical record with error $\leq \varepsilon$ (mis-record probability) over time T (record retention time). We treat (ε, T) as externally fixed operational tolerances defining what counts as a stable record; isosymmetry is defined relative to a chosen tolerance regime.

Tolerance robustness. Although Definition 4.1(iii) introduces tolerance parameters (ε, T), the resulting isosymmetric classification is stable above a minimal record-formation threshold. Below this threshold, no system supports fact creation and isosymmetry is undefined. Above it, variations in (ε, T) do not change the admissible task structure or its compositional rules. In this sense, isosymmetry is robust, not finely tuned: it is defined relative to the *existence* of stable records, not to arbitrary precision choices.

This correspondence need not preserve microscopic degrees of freedom, spatial embedding, or dynamical details; it preserves only the structural invariants (i)–(iii).

Definition 4.2 (Operational Isosymmetry). Two physical systems *A* and *B* are *operationally isosymmetric* if, for every finite discrimination, composition, and correlation task implementable by admissible instruments on *A* with a given resource scaling (distinguishable states, committed outcomes, exported correlations), there exists a corresponding task on *B* with identical scaling behaviour, and vice versa.

Isosymmetry is therefore an equivalence relation defined over operational task structure and resource scaling, not over microscopic constitution or dynamics.

4.3 Conditions for Isosymmetry

Isosymmetry requires equivalence in three respects:

1. **Equivalence of admissible state space:** The sets of physically realisable global states must be structurally identical with respect to distinguishability and correlation.
2. **Equivalence of composition:** Admissible joint configurations must combine according to the same constraint rules, yielding identical compositional structure.
3. **Equivalence of resolution:** At critical thresholds, admissible alternatives must collapse into committed outcomes in the same structural manner.

These conditions together define an *isosymmetric class*. Membership in such a class guarantees that systems share the same abstract quantum structure, regardless of their physical substrate.

4.4 What Isosymmetry Does Not Require

Systems need not share the same Hamiltonian, the same interaction mechanisms, or the same spacetime description. They need not be reducible to one another, nor derivable from a common microscopic theory. Isosymmetry is indifferent to material constitution; it is sensitive only to admissible constraint resolution.

4.5 Why the Same Task Class?

Isosymmetry explains why systems that realise the same admissible task class share quantum structure. It does not, by itself, explain *why* photons, electrons, atoms, and superconducting circuits all realise the same class. Three non-exclusive possibilities exist:

Contingent universality. It may be a deep empirical fact about our universe that all stable fact-producing substrates satisfy the same admissibility constraints. On this view, isosymmetry explains structure given the fact, but not the fact itself.

Selection effects. Only task classes supporting global consistency, stable records, and compositional closure permit complex observers and long-lived experiments. Non-isosymmetric classes may exist in principle but be unobservable.

Constraint uniqueness. The admissibility conditions themselves may be sufficiently restrictive that only one nontrivial task class exists. Spekkens's toy model [15] demonstrates that operational similarity alone does not guarantee isosymmetry: the model reproduces many quantum features but lacks entanglement and is not isosymmetric to quantum systems. This illustrates that isosymmetric equivalence requires the full admissibility package, not merely superficial structural resemblance. Exploring whether admissibility uniquely determines the quantum task class is left to future work.

The present paper remains neutral between these interpretations while providing the structural explanation common to all three.

---

5. Consequences of Isosymmetry

The principle of isosymmetry has immediate and far-reaching consequences. If distinct physical systems belong to the same isosymmetric constraint class, then the abstract structure used to describe them is not a modelling choice but a structural necessity.

5.1 Universality of Hilbert Space Structure

Hilbert space appears in quantum mechanics as a universal representational framework. States are vectors, observables are operators, and probabilities are computed via inner products. In standard formulations, this structure is postulated rather than derived.

Under isosymmetry, Hilbert space arises as the minimal structure capable of representing admissible global constraint states with graded distinguishability and superposition. If two systems admit the same constraint resolution behaviour, then their state spaces must support identical notions of overlap, orthogonality, and combination.

Lemma 5.1 (Admissibility ⇒ convex operational structure). Under Definition 3.1, the set of preparation procedures induces a convex state space $\mathcal{S}$ (mixtures correspond to randomised preparations), and readouts induce affine functionals ("effects") e: $\mathcal{S} \to [0,1]$. Transformations act as affine maps on $\mathcal{S}$.

Justification. Randomisation is a physical operation available within Instr(*D*); irreversible records imply well-defined outcome probabilities; resource bounds restrict to finite scenarios. The convex structure is therefore not assumed but constructed from operational primitives.

Lemma 5.2 (Explicit Mapping from Admissibility to Reconstruction Assumptions). Let Instr(*D*) satisfy Definition 3.1 and let admissible experiments be closed under sequential and probabilistic composition. Then the operational structure induced by Instr(*D*) satisfies the following properties:

(P1) Convexity of states (mixtures). For any two admissible preparations P_1, P_2 and any $\lambda \in [0,1]$, there exists an admissible preparation P_λ that implements P_1 with probability λ and P_2 with probability 1−λ. The resulting equivalence classes of preparations define a convex state space $\mathcal{S}$.

(P2) Affinity of measurement statistics. For any admissible measurement M with outcomes o, the map $s \mapsto p(o|s,M)$ is affine on $\mathcal{S}$.

(P3) Existence of reversible transformations. Any admissible reversible operation T ∈ Instr(*D*) acts as an affine automorphism of $\mathcal{S}$. The set of such transformations forms a group G acting on $\mathcal{S}$.

(P4) Finite information capacity. Finite distinguishability (Definition 3.1(i)) implies: for each system there exists a finite N such that no admissible measurement perfectly distinguishes more than N states. This is the operational "capacity" parameter used in reconstructions.

(P5) Composition and non-signalling. For spacelike-separated instruments A and B, admissible joint experiments define a composite state space $\mathcal{S}$_AB whose marginal statistics are independent of remote measurement choice (non-signalling). Equivalently: reduced states exist as well-defined marginals.

(P6) Continuous reversibility (additional regularity). If, in addition, admissible reversible operations vary continuously with control parameters and act transitively on pure states within a connected component, then the "continuous reversibility" assumption used by Hardy and Masanes–Müller holds.

(P7) Local tomography (explicit assumption, when used). If, in addition, joint states are determined by statistics of local measurements (tomographic locality), then the "local tomography" assumption used in several reconstructions holds.

(P8) Purification (explicit assumption, only for CDP-type routes). If, in addition, every mixed state arises as a marginal of a pure state on a larger system, then the "purification" assumption holds (Chiribella–D'Ariano–Perinotti).

Proof. (P1)–(P3) follow from closure under probabilistic composition and the existence of irreversible records that define stable outcome frequencies. (P4) is exactly Definition 3.1(i). (P5) is the operational content of §3.3: admissibility is global, but local instrument choices cannot change remote marginals. (P6)–(P8) are stated explicitly as additional regularity/representation assumptions and are not claimed to follow from Definition 3.1 alone. ■

Derivation of (P1)–(P2). Implement P_λ by coupling the system to a classical randomiser that irreversibly records a bit b with $p(b=1) = \lambda$ and triggers P_1 if b=1 else P_2. By construction, for any measurement outcome o,

$$p(o|P\_\lambda,M) = \lambda\ p(o|P_1,M) + (1-\lambda)\ p(o|P_2,M),$$

so the operational equivalence class s_λ satisfies $s\_\lambda = \lambda s_1 + (1-\lambda)s_2$ and measurement functionals are affine.

Derivation of (P5) (non-signalling). If local choice y at B could change marginal statistics at A, then B could transmit a controllable bit to A by switching y, and A could irreversibly record it. This would contradict the operational separability implicit in admissible instrument composition (and would empirically violate relativistic causal constraints). Hence admissible joint experiments must satisfy $p(a|x,y) = p(a|x)$ for all x, y, so reduced states exist as well-defined marginals.

Reconstruction ingredient	Where it appears	Status here
Probabilistic mixtures / convexity	Hardy, Masanes–Müller, CDP (implicit)	Derived (P1–P2)
Reversible transformations as symmetries	all	Derived (P3)
Finite capacity / distinguishability	Hardy ("information capacity")	Derived (P4)
Non-signalling / well-defined marginals	all	Derived (P5)
Continuous reversibility	Hardy, Masanes–Müller	Additional (P6)
Local tomography	Hardy-style routes, many GPT→QM theorems	Additional (P7)
Purification	CDP	Additional (P8)
Hardy "simplicity"	Hardy	Replaced by isosymmetry-class equivalence

Remark. Hardy's "simplicity" axiom is not entailed by admissibility; it is replaced by isosymmetry-class identification (same capacity + same task structure ⇒ same representation class). Purification is a stronger assumption about how mixed states arise; it remains optional and is required only for CDP-type reconstruction routes. Continuous reversibility follows naturally in many physical systems from controllable Hamiltonian evolution. Local tomography is a representation assumption, not a physical postulate derived from admissibility.

Theorem 5.1 (Hilbert Space Representation). Lemma 5.2 derives (P1)–(P5) explicitly from admissibility, and states (P6)–(P8) as optional regularity assumptions aligned with specific reconstruction routes. Under these assumptions, reconstruction theorems imply a Hilbert-space representation; companion work [10] selects ℂ as the amplitude field. Specifically, any formalism capable of faithfully encoding admissible global constraint states with:

- (i) graded distinguishability (states may be partially but not fully discriminable),
- (ii) reversible composition (joint states can be formed and factored without loss), and
- (iii) non-signalling extension (local operations on one subsystem do not alter statistics of another)

must be equivalent to a complex Hilbert space description up to operational isomorphism.

Here 'equivalent' means: there exists a mapping of the theory's states and effects into the projective space and POVM effects of a complex Hilbert space that preserves outcome probabilities for all admissible experiments.

Clarification. Theorem 5.1 is not a reconstruction of quantum theory from axioms. It is a *representation theorem*: it states that any formalism capable of faithfully encoding isosymmetric admissible constraint classes must collapse to Hilbert space structure. The theorem does not derive quantum mechanics; it shows that no alternative representation of the same constraint class exists.

Proof sketch. The conditions (i)–(iii) correspond to the axioms from which reconstruction theorems derive Hilbert space structure [5–8]. The deeper question is why *complex* Hilbert space specifically. This is established in companion work [10] via Galois-theoretic selection: requiring that observable predictions be invariant under automorphisms of the amplitude field, combined with interference and isotropy requirements, uniquely selects ℂ. Real numbers are excluded because $\mathrm{Aut}(\mathbb{R}/\mathbb{R})$ is trivial, precluding continuous phase. Quaternions are excluded because $\mathrm{Aut}(\mathbb{H}/\mathbb{R}) \cong SO(3)$ conflicts with permutation symmetry of distinguishable configurations—joint invariance destroys phase-sensitivity. Complex numbers, with $\mathrm{Aut}(\mathbb{C}/\mathbb{R}) = \{\text{identity}, \text{conjugation}\}$, satisfy all constraints. The resulting state space is necessarily a complex Hilbert space. ■

This does not imply that physical systems are literally vectors in an abstract space. Rather, Hilbert space functions as a canonical encoding of admissible constraint relations. Its universality follows from the universality of admissibility, not from the material properties of the systems it represents.

Remark. The derivation of Hilbert space structure from information-geometric principles, including the emergence of the inner product from Fisher-Rao geometry and phase structure, is developed in detail in companion work [10].

5.2 Universality of Tensor-Product Composition

The tensor product rule governs how composite quantum systems are represented. When two systems are combined, their joint state space is not a Cartesian product but a tensor product, allowing for states that cannot be decomposed into independent components.

From the perspective of isosymmetry, this compositional rule is unavoidable. If admissibility applies to global configurations rather than local subsystems, then the representation of joint systems must preserve the possibility of global constraints. The tensor product is precisely the structure that encodes such global admissibility while retaining compatibility with local descriptions.

The tensor product is not the unique mathematically consistent composition rule for operational theories; GPTs admit alternatives including minimal and maximal tensor products [13]. Isosymmetry selects the quantum tensor product because alternatives either permit signalling or fail to support stable record formation under Bell-type correlations [14].

Because isosymmetry equates systems at the level of admissible composition, any two isosymmetric systems must combine according to the same tensor rules. The tensor product is therefore not a peculiarity of wave mechanics or a postulate of quantum kinematics, but an invariant of isosymmetric constraint classes.

5.3 Universality of Entanglement

Entanglement is often treated as a distinctive feature of quantum mechanics requiring special explanation. Within the present framework, it emerges naturally.

Proposition 5.2. If a system belongs to an isosymmetric class that admits global constraint resolution with non-signalling correlations, and if the admissible task set includes Bell-type experiments whose observed correlations violate a Bell inequality while remaining non-signalling, then non-factorizable (entangled) states are structurally required.

Proof. Suppose admissibility requires global consistency of fact creation: outcomes at separated components must not contradict one another, yet local operations must not signal. A representational framework that forbids entanglement would restrict joint states to product form, but product states cannot violate Bell inequalities. If the admissible task set empirically includes Bell-violating correlations (as observed in quantum systems), then any adequate representation must include non-factorizable states. (Note: this is an empirical premise—that Bell violations occur—not a consequence of Definition 3.1 alone.) ■

Entangled states are simply admissible global constraint states that cannot be factorised into independent local configurations. If a system belongs to an isosymmetric class that admits global constraint resolution, then entanglement is not optional—it is required.

This explains why entanglement appears universally across quantum systems, independent of scale or physical implementation. Photons, atoms, spins, and superconducting circuits all support entanglement because they belong to the same isosymmetric class. Entanglement is not a dynamical effect added to otherwise separable systems; it is a structural feature of admissibility itself.

Remark on probability. Isosymmetry constrains which outcome structures are admissible, not how probability weight is assigned among them. The Born rule, which determines measurement probabilities, is derived separately from the structure of admissible resolution [9]. Isosymmetry explains why all quantum systems share the same *framework* for probability; it does not determine the probabilities themselves.

5.4 Worked Example: Photon Polarisation and Electron Spin

Consider two paradigmatic quantum systems: photon polarisation and electron spin-1/2.

Photon polarisation:

- Substrate: Electromagnetic field mode
- Degrees of freedom: Polarisation state (horizontal, vertical, or superposition)
- Physical mechanism: Maxwell's equations, optical elements

Electron spin:

- Substrate: Fermionic matter field
- Degrees of freedom: Spin state (up, down, or superposition along any axis)
- Physical mechanism: Dirac equation, magnetic fields

These systems differ in every microscopic respect: one involves bosonic fields, the other fermionic particles; one couples to electric fields, the other to magnetic fields; one propagates at light speed, the other can be stationary.

Yet both belong to the same isosymmetric class because:

1. Both support exactly 2 perfectly distinguishable states (Definition 3.1(i))
2. Both admit continuous reversible transformations between pure states (SU(2) structure)
3. Both compose via tensor products yielding 4-dimensional joint spaces with entangled states
4. Both exhibit identical correlation structure (same Bell inequality violations)
5. Both undergo irreversible resolution producing identical statistical patterns (Born rule)

The identical quantum structure is not a coincidence of similar dynamics—the dynamics are entirely different. It is a consequence of isosymmetry: both systems satisfy identical admissibility constraints and therefore belong to the same constraint class.

5.5 Worked Example: Why Classical Phase Space Fails Isosymmetry

Classical mechanics is often thought of as the "simplest" physical theory. If isosymmetry merely redescribed universality, one might expect classical phase-space systems to belong to the same isosymmetric class as quantum systems. This subsection shows explicitly that they do not.

The failure is not dynamical but structural: classical phase space violates the finite distinguishability requirement of admissibility (Definition 3.1(i)), and therefore realises a different task class.

Setup. Consider a classical particle with phase space $\Gamma = \mathbb{R}^{2n}$, with points (q, p) representing exact position and momentum. Preparations correspond to probability densities ρ(q, p) on Γ; measurements correspond to partitions of phase space into disjoint measurable regions; dynamics are Hamiltonian flows preserving phase-space volume. This defines a perfectly well-formed classical operational theory.

Step 1: Continuous distinguishability. For any finite region $R \subset \Gamma$, classical mechanics permits arbitrarily fine partitions into N disjoint subregions, with N unbounded. For any N, there exists a measurement that perfectly distinguishes which subregion the system occupies. Hence classical phase space admits arbitrarily many perfectly distinguishable states in a single measurement—there is no upper bound N_max.

Step 2: Violation of finite capacity. Definition 3.1(i) requires that at most n states can be perfectly discriminated in a single measurement. Classical phase space violates this requirement fundamentally: no finite information capacity exists, distinguishability scales without bound as measurement resolution increases, and there is no natural cutoff enforcing a maximal set of mutually exclusive outcomes. This is not a practical limitation but a structural property of classical theory.

Step 3: Consequences for task structure. Because finite distinguishability fails:

- No finite-dimensional convex state space exists (the space of perfectly distinguishable preparations is infinite-dimensional even locally)
- No finite capacity parameter can be defined (Hardy's "information capacity" is undefined)
- No admissible global resolution threshold exists (there is no point at which underdetermined alternatives are forced into a finite outcome set)
- No nontrivial isosymmetry class is shared with quantum systems

Step 4: Contrast with quantum systems.

Property	Classical phase space	Quantum system
Perfectly distinguishable states	Unbounded	Finite ($\leq$ d)
Capacity parameter	None	Well-defined
Global resolution thresholds	Absent	Present
Isosymmetric with quantum systems	✗	✓

Interpretation. Classical mechanics fails isosymmetry not because it is "less rich" but because it is *too permissive*: it allows arbitrarily fine discrimination, unlimited information extraction in a single measurement, and no intrinsic threshold for irreversible fact creation. Quantum theory sits at a critical boundary—continuous structure (phases, superpositions) but finite distinguishability and bounded capacity. This balance is exactly what admissibility encodes.

This example demonstrates that isosymmetry is not a restatement of universality: it excludes classical mechanics for principled reasons and identifies finite distinguishability as the structural hinge separating quantum from classical constraint classes.

---

6. Relation to Criticality

The principle of isosymmetry explains why quantum structure is universal across physical realisations, but it does not, by itself, explain when or how irreversible outcomes occur. That role is played by *criticality*. It is essential to distinguish clearly between the explanatory domains of these two concepts.

Principle 6.0 (Threshold Invariance). Crossing a critical threshold changes which admissible alternatives remain (resolution selects among them), but not the underlying admissible task structure governing composition and correlation. Criticality is a selection on states *within* a fixed constraint class, not a transition *between* constraint classes.

On this view, criticality is to isosymmetry as symmetry breaking is to symmetry: the dynamics selects a branch, but the kinematic arena remains fixed.

6.1 What Criticality Explains

Criticality concerns the conditions under which a physical system transitions from reversible evolution to irreversible commitment. Below critical thresholds, systems may occupy admissible states that remain underdetermined, oscillatory, or globally constrained without producing definite outcomes. At critical thresholds, admissibility constraints force resolution: one alternative is selected, entropy increases, and a fact is created. This transition is structural rather than dynamical, reflecting limits on capacity, consistency, and irreversible record formation.

6.2 Complementary Explanatory Domains

Isosymmetry does not determine the location of critical thresholds or the dynamics leading up to them. Instead, it classifies systems according to the structure of admissible resolution once such thresholds are reached. Two systems may differ in scale, interaction strength, or environmental coupling, and therefore reach criticality under different conditions, while still belonging to the same isosymmetric class.

Question	Answered by
Why do different systems have the same quantum structure?	Isosymmetry
When does irreversible resolution occur?	Criticality
What determines measurement outcomes?	Resolution dynamics
Why are outcomes probabilistic?	Born rule (derived separately)

Criticality answers the question of *when* irreversible resolution occurs, while isosymmetry answers the question of *why* the same resolution structure applies across systems. The two principles are complementary rather than competing.

6.3 Entanglement and Global Resolution

This distinction is particularly important for understanding entangled systems. In an entangled configuration, admissibility applies to the global constraint state rather than to local subsystems. When criticality is reached, resolution occurs at the level of the global constraint, producing correlated outcomes without the need for signal propagation or causal influence between spatially separated components. Isosymmetry ensures that this pattern of global resolution is identical across physical implementations, while criticality determines when resolution is enforced.

---

7. Comparison with Other Universality Notions

Universality is a familiar concept in physics, appearing in several well-established contexts. It is important to clarify how isosymmetry differs from, and complements, existing notions.

7.1 Renormalisation Group Universality

In critical phenomena, systems with different microscopic interactions flow toward the same fixed points and share identical scaling behaviour. However, renormalisation-group universality concerns robustness under coarse-graining and scale transformation. It presupposes an underlying theory and does not explain why the same abstract kinematic and compositional structure applies to all quantum systems at all scales. Isosymmetry operates prior to dynamics and scale, classifying systems by admissible resolution structure rather than by flow toward fixed points.

7.2 Effective Field Theory

Effective field theory explains why low-energy observables are largely independent of high-energy physics. While effective theories successfully capture emergent behaviour within a given kinematic framework, they presuppose quantum structure. They do not explain why the same Hilbert space formalism, tensor composition, and entanglement rules apply across disparate physical domains.

7.3 Decoherence

Decoherence-based accounts explain why certain bases become preferred under environmental interaction and why interference terms are suppressed in practice. However, decoherence assumes the existence of entanglement, tensor-product composition, and Hilbert space from the outset. It therefore cannot ground their universality. Isosymmetry explains why these structures must be present before any decoherence process can occur.

7.4 Information-Theoretic Reconstructions

Reconstruction programs derive quantum formalism from operational or informational axioms [5–8]. These approaches successfully derive aspects of the formalism, demonstrating that quantum structure follows from constraints such as tomographic locality, continuous reversibility, and purification.

Isosymmetry operates at a complementary level. Reconstruction programs answer: "Given that systems satisfy these axioms, what structure follows?" Isosymmetry answers: "Why do all fact-producing systems satisfy axioms of this type in the first place?"

This paper establishes the substrate-independence principle (isosymmetry); the amplitude-field selection and Hilbert-space emergence are proven constructively in [10] and used here as the structural backbone.

The full derivation chain is (reconstruction axioms appear here not as derived results, but as formal codifications of the admissible task structure identified by isosymmetry):

Admissibility constraints (this paper)
↓
Isosymmetric constraint classes
↓
Operational axioms (reconstruction programs)
↓
Galois invariance + isotropy [10]
↓
Complex Hilbert space (uniquely)

Specifically, companion work [10] completes the chain by proving that among the three candidate amplitude fields permitted by Frobenius's theorem ($\mathbb{R}$, $\mathbb{C}$, $\mathbb{H}$), only complex numbers satisfy the joint requirements of interference, Galois invariance, and isotropy. This transforms the "why complex?" question from a brute fact into a theorem.

Isosymmetry thus provides a physical grounding for the axioms that reconstruction programs take as starting points, while [10] explains why those axioms force complex rather than real or quaternionic structure.

Multiple reconstruction routes—operational [5–7], diagrammatic [16], and information-geometric [10]—converge on the same structure despite differing formalisms. This convergence is itself evidence for isosymmetry: the routes identify the same constraint class through different lenses.

7.5 Structural Realism

The claim that quantum structure reflects equivalence of constraint classes rather than shared material ontology resonates with structural realism in philosophy of physics [11]. Structural realists hold that what persists across theory change and across physical instantiation is structure rather than objects. Isosymmetry provides a precise criterion for structural equivalence: two systems are structurally identical (in the relevant sense) if and only if they are isosymmetric.

---

8. Implications and Limits

8.1 Implications

Principled explanation for substrate-independence. The identical kinematic and compositional structure observed across quantum systems is not an unexplained coincidence or a matter of modelling convenience. It reflects a deeper equivalence at the level of admissible constraint resolution.

Shift from dynamics to admissibility. Attempts to derive quantum theory purely from underlying dynamics or specific interaction mechanisms may be misdirected. Quantum structure is not imposed by particular forces or constituents, but by the requirements of admissible global resolution.

Boundary conditions for quantum applicability. Systems that fail to meet admissibility requirements—for example, systems lacking sufficient capacity for global constraint formation or irreversible resolution—would not belong to a quantum isosymmetric class. This opens the possibility of identifying regimes in which quantum structure breaks down, not because of new dynamics, but because admissibility conditions are violated.

8.2 Limits

The present account has deliberate limits:

- Isosymmetry does not determine the probabilities of specific outcomes (this requires the Born rule, derived separately).

- Isosymmetry does not specify the dynamics by which systems approach criticality.
- Isosymmetry does not resolve the selection of particular measurement results.
- Isosymmetry does not replace existing quantum dynamics.

These questions lie outside the scope of the present paper and are addressed by complementary principles.

8.3 Falsifiability

Isosymmetry makes the following empirically testable predictions:

Prediction. Any physical system capable of:

- (i) supporting stable, irreversible records, and
- (ii) exhibiting at least two distinguishable states with continuous reversible transformations between them

will exhibit quantum structure (Hilbert space, tensor composition, entanglement).

Isosymmetry would be *falsified* by either of the following:

A. Correlation falsifier (strongest). A fact-producing system satisfying Definition 3.1 that exhibits correlations outside the quantum set (e.g., above Tsirelson bounds) in a controlled Bell scenario, while maintaining stable classical records of outcomes.

B. Composition falsifier (cleanest). A fact-producing system with two-level controllable subsystems whose composite does not admit tensor-product structure or entangled extensions—i.e., only product composites are physically realisable despite both subsystems individually satisfying admissibility.

Candidate domains for falsification:

- Exotic post-quantum GPT proposals with novel composition rules
- Indefinite causal order scenarios, if they produced superquantum correlations with stable records
- Any verified violation of information causality with stable classical storage
- Gravitational or Planck-scale regimes where admissibility conditions might fail

Remark on superquantum theories. Generalised probabilistic theories (GPTs) such as PR-box theories formally permit stronger-than-quantum correlations [17]. These serve as theoretical counterexamples that illuminate the boundaries of isosymmetric classes. PR-box correlations would trivialise communication complexity [18], collapsing the resource scaling that defines admissibility (Definition 3.1(iii)). The quantum correlation boundary, precisely characterised by Navascués et al. [19], may therefore reflect the limit of correlations compatible with non-trivial fact production. Crucially, such theories are widely conjectured to be inadmissible under stable record formation and information causality constraints [12]. The non-existence of physical PRboxes would thus reflect not merely empirical accident but the deeper constraint that admissibility imposes on correlation structure.

To date, no physical system satisfying the falsification criteria has been observed.

---

9. Conclusion

The universality of quantum structure across physical realisations has long been treated as a brute fact or convenient assumption. We have argued that it admits a principled explanation.

Isosymmetry—equivalence at the level of admissible constraint resolution—provides the missing principle. Systems that support the same patterns of finite distinguishability, reversible composition, and irreversible commitment necessarily exhibit identical quantum structure. Hilbert space, tensor products, and entanglement are not imposed separately on each physical domain; they are invariants of isosymmetric constraint classes.

This paper is part of a reconstruction programme that derives quantum mechanics from distinguishability principles:

Paper	Question Answered	Method
This paper	Why is quantum structure substrate-independent?	Isosymmetry over constraint classes
Hilbert space paper [10]	Why complex numbers specifically?	Galois invariance excludes $\mathbb{R}$ and $\mathbb{H}$
Born rule paper [9]	Why $P = \|\psi\|^2$?	Coarse-graining geometry

Together, these results transform quantum universality from mystery to theorem: the kinematic structure (complex Hilbert space), the compositional structure (tensor products), and the probability rule (Born rule) all follow from the geometry of distinguishable configurations and the requirements of consistent fact production.

This perspective relocates the source of quantum universality from material constitution to structural admissibility. The question "why does quantum mechanics apply to photons and electrons alike?" receives a precise answer: because photons and electrons are isosymmetric—they belong to the same constraint class, and constraint class determines structure.

Isosymmetry thus completes a chain of explanation: reconstruction programs show that certain axioms entail quantum formalism; isosymmetry shows why fact-producing systems satisfy those axioms. Together, they transform quantum universality from mystery to theorem.

References

[1] C. Rovelli, "Relational quantum mechanics," *International Journal of Theoretical Physics* **35**, 1637 (1996).

[2] M. Van Raamsdonk, "Building up spacetime with quantum entanglement," *General Relativity and Gravitation* **42**, 2323 (2010).

[3] C. Cao, S. M. Carroll, and S. Michalakis, "Space from Hilbert space: Recovering geometry from bulk entanglement," *Physical Review D* **95**, 024031 (2017).

[4] N. Huggett and C. Wüthrich, "Out of Nowhere: The emergence of spacetime in quantum theories of gravity," arXiv:2101.03321 (2021).

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

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

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

[8] B. Dakić and Č. Brukner, "Quantum theory and beyond: Is entanglement special?" in *Deep Beauty*, ed. H. Halvorson (Cambridge University Press, 2011).

[9] K. Taylor, "The Born rule from coarse-graining geometry," VERSF Working Paper (2025). [Derives $P = |\psi|^2$ from distinguishability geometry without assuming probability structure.]

[10] K. Taylor, "Complex Hilbert space from distinguishability principles," VERSF Working Paper (2025). [Derives complex Hilbert space from Galois invariance on distinguishability geometry; proves $\mathbb{R}$ excluded (no continuous phase), $\mathbb{H}$ excluded (SO(3) automorphisms clash with isotropy), $\mathbb{C}$ uniquely selected. Available as supplementary manuscript.]

[11] J. Ladyman, "What is structural realism?" *Studies in History and Philosophy of Science* **29**, 409 (1998).

[12] M. Pawłowski et al., "Information causality as a physical principle," *Nature* **461**, 1101 (2009).

[13] J. Barrett, "Information processing in generalized probabilistic theories," *Physical Review A* **75**, 032304 (2007).

[14] H. Barnum and A. Wilce, "Post-classical probability theory," in *Quantum Theory: Informational Foundations and Foils*, eds. G. Chiribella and R. W. Spekkens (Springer, 2016).

[15] R. W. Spekkens, "Evidence for the epistemic view of quantum states: A toy theory," *Physical Review A* **75**, 032110 (2007).

[16] J. H. Selby, C. M. Scandolo, and B. Coecke, "Reconstructing quantum theory from diagrammatic postulates," *Quantum* **5**, 445 (2021).

[17] S. Popescu and D. Rohrlich, "Quantum nonlocality as an axiom," *Foundations of Physics* **24**, 379 (1994).

[18] W. van Dam, "Implausible consequences of superstrong nonlocality," *Natural Computing* **12**, 9 (2013).

[19] M. Navascués, S. Pironio, and A. Acín, "A convergent hierarchy of semidefinite programs characterizing the set of quantum correlations," *New Journal of Physics* **10**, 073013 (2008).

Appendix A: Summary of Amplitude Field Selection

For completeness, we summarise the Galois-theoretic argument establishing that complex numbers are the unique viable amplitude field, proven constructively in companion work [10].

Setup. Let Λ be a finite set of distinguishable configurations with $|\Lambda| \geq 3$. Amplitudes are functions $\psi: \Lambda \to F$ for some division algebra F over $\mathbb{R}$. By Frobenius's theorem, the candidates are $F \in \{\mathbb{R}, \mathbb{C}, \mathbb{H}\}$.

Constraints imposed:

1. *Interference:* F must support continuous phase structure
2. *Galois invariance:* Observable predictions must be invariant under $\mathrm{Aut}(F/\mathbb{R})$
3. *Isotropy:* Predictions must be invariant under permutations of Λ (when configurations are symmetrically distinguishable)
4. *Regularity (Taylor Limit):* Only continuous automorphisms are physically admissible

Regularity (Taylor Limit) — Physical Justification. The Taylor Limit requires that observable probabilities vary smoothly under small perturbations of admissible preparations and transformations. Physically, this reflects:

- finite experimental resolution (no infinitely precise measurements),
- the impossibility of infinitely sharp control, and
- the empirical success of perturbative modelling in physics.

Without such regularity, pathological representations (e.g., discontinuous automorphisms of $\mathbb{C}$ constructed via Zorn's lemma) become formally admissible but are physically meaningless—

they correspond to nowhere-measurable functions incompatible with any laboratory procedure. The Taylor Limit therefore excludes mathematical artefacts rather than physical possibilities.

Results:

Field	Aut(F/ℝ)	Verdict	Reason
ℝ	{id}	Excluded	Trivial automorphism group cannot support continuous phase; only discrete (±1) phase available
ℍ	SO(3)	Excluded	Automorphism group too large; SO(3) acting on imaginary quaternions conflicts with S_n isotropy—joint invariance destroys phase-sensitivity
ℂ	{id, conj}	Selected	Two-element group supports interference while commuting with all permutations

Theorem (Selection of ℂ). Under constraints 1–4, complex numbers are the unique amplitude field compatible with interference, Galois invariance, and isotropy. The resulting state space is necessarily a complex Hilbert space.