

From Overdetermination to Inevitability

A Unification of the Five VERSF Routes to Quantum Probability

Keith Taylor *VERSF Theoretical Physics Programme*

Plain-Language Summary

Of all the rules in quantum mechanics, the rule for *probability* is the strangest. It says that if a quantum system is described by an amplitude with components c_1, c_2, \dots , then the probability of obtaining outcome i upon measurement is $|c_i|^2$ — the *square* of the magnitude of the component. This is the **Born rule**, after Max Born who proposed it in 1926.

The rule works perfectly. Every quantum experiment ever performed has confirmed it. But for nearly a century, physicists have been unable to agree on *why* the probability should be a square rather than, say, a fourth power, an absolute value, or a logarithm. Born had no derivation; he simply postulated the rule. Most quantum mechanics textbooks still simply postulate it. The major attempts to derive it from deeper principles — Gleason in 1957, Deutsch–Wallace in the 1990s, Zurek in the 2000s, Hardy in 2001, Masanes–Müller in 2011 — each succeed mathematically, but each rests on an assumption no less strange than the rule it derives. And every one of them quietly takes for granted the underlying mathematical stage on which the rule is supposed to act: a *complex Hilbert space* of quantum states. They derive the rule but presuppose the stage. The prior question — why does the universe use *that* particular mathematical structure in the first place? — is left untouched.

Earlier work in the VERSF programme established something stronger about the rule itself. **Five completely independent derivations within the framework all converge on the Born rule.** They share no proof step, no premise, no mathematical machinery beyond the framework's three foundational principles. They begin from completely different physical considerations — the structure of reversible dynamics, the consistency of measurement records, the geometry of paths through configuration space, the direct elimination of alternative probability rules, and the maximum-entropy structure of measurement processes — and they all terminate at the same answer: $p_i = |c_i|^2$.

This paper began with one question: *is the convergence accidental?* Is it a coincidence that five logically independent routes converge on the same rule, or is there a single underlying structure that the five routes are each expressing in their own vocabulary?

We prove the latter is the case. The five routes are not five independent confirmations. They are **five projections of a single object** — what we call the **admissibility framework**, the structure of constraints any fact-producing physical theory must satisfy. Each route makes one face of that framework operationally explicit. They converge because they are all looking at the same thing from different angles, and the Born rule is the unique probability rule the framework admits.

But in working this out, the paper grew larger than we initially planned. The same admissibility framework that forces the Born rule turns out to also force the *kinematic stage on which the rule plays out* — the very thing earlier reconstructions had to take for granted. Three substrate results within the proof establish that the framework selects:

- **complex-valued amplitudes** rather than real or quaternionic ones — only complex numbers give a consistent rule for combining systems into larger systems without producing hidden global parameters or inconsistent marginal frequencies;
- **the Hilbert-space inner product** — forced by the requirement that the symmetry group of admissible reversible transformations be compact, which in turn is forced by the bound on operationally distinguishable states in any bounded system;
- **unitary evolution** of pre-measurement dynamics — forced by the requirement that distinctions between physical states be preserved during evolution (otherwise commitment leaks into the dynamics, or new facts appear without any commitment event having occurred).

In other words: the same conditions that force the Born rule also force the entire mathematical machinery of quantum mechanics. The earlier reconstructions had to assume Hilbert space and then derive the probability rule. Here, both are derived together from the same minimal starting points.

The strongest defensible claim of the paper is therefore not just about the Born rule:

Quantum kinematics (complex Hilbert space, unitary dynamics) and the Born rule jointly emerge as the unique admissible structure consistent with finite distinguishability, irreversible fact-formation, observer-invariance, compositional consistency, and reversible pre-commitment dynamics.

Stated less technically: any physical theory that produces *facts* (rather than merely possibilities), in which bounded systems have finitely many distinguishable states, in which systems can be combined into larger systems whose marginal statistics are observer-invariant, and whose internal dynamics preserve distinctions between states — must have complex-Hilbert-space structure and Born-rule probability. There is no other admissible architecture.

This changes what quantum mechanics *is*. It is not one theoretical framework among several plausible alternatives that the universe happens to instantiate. It is the unique architecture compatible with the conditions for a fact-producing, composable, observer-invariant physical theory. The Born rule, in particular, is no longer a postulate accepted because it works empirically — it is the only probability rule the architecture can have, on the only kinematic stage the architecture admits.

In one sentence: **the Born rule, complex Hilbert space, and unitary dynamics are jointly forced by the conditions for a fact-producing, composable, observer-invariant physical theory.**

Abstract

Earlier work within the Void Energy-Regulated Space Framework (VERSF) established that five operationally independent derivations converge exactly on the Born rule, constituting structural overdetermination of unprecedented strength in the foundations of quantum mechanics. The present paper goes further. We show that the five derivations are not merely consistent but are *structurally equivalent projections of a single admissibility framework* grounded in the VERSF kernel principles of finite distinguishability, irreversible commitment, and observer invariance, together with derived structural conditions of compositional consistency and reversible pre-commitment dynamics.

We prove the **Admissible Probability Uniqueness Theorem**: any probability assignment compatible with the admissibility framework $\mathfrak{A} = A0 \wedge A1 \wedge A2 \wedge B1 \wedge B2$ must be quadratic in amplitudes, with the unique form $p_i = |c_i|^2$. No alternative rule satisfies the joint constraints. The Born rule is therefore not only overdetermined but **the unique admissible probability measure within any fact-producing framework satisfying \mathfrak{A}** .

The proof strategy is **unifying rather than novel**: we do not invent new mathematics. We identify seven lemmas whose joint satisfaction the theorem requires. The first three are *substrate lemmas* that derive Hilbert structure from \mathfrak{A} itself: a $U(1)$ phase substrate from $A1 \wedge A2$; local tomography from $A0 \wedge A1 \wedge B1$, selecting \mathbb{C} as the amplitude field over \mathbb{R} and \mathbb{H} ; and compactness of the reversible-dynamics group from $A1 \wedge B2$, forcing an invariant inner product. The remaining four are *rule lemmas* — the existing five Routes restated as constraints on the derived substrate — which fix the probability rule. **No external Hilbert-space presupposition remains in the proof.** Each constraint, taken alone, leaves room for alternative probability structures along directions the other constraints exclude; only joint satisfaction forces the Born rule. This explains both *why* the five routes converge (they constrain the same object) and *why* the convergence is forced (no slack remains in the admissibility structure).

We address head-on the most pointed referee objection — that the framework smuggles the Born rule into its admissibility conditions — by showing that the conditions specify only what counts as a fact, what persists, how systems compose, and what observer-invariance requires. **None of these conditions mentions probability.** Probability emerges as the unique measure consistent with the structure they jointly impose. The analogy is with the entropy function in thermodynamics: it is not built into the laws of thermodynamics, but is forced once those laws are imposed.

The result reframes the Born rule from a postulate or convergent outcome into a **structural necessity**: the unique admissible measure compatible with the minimal conditions required for physical reality to exist, persist, and be observed.

Table of Contents

Front matter

- Plain-Language Summary
- Abstract
- Table of Contents
- Notation Conventions

1. Introduction: From Convergence to Necessity

- 1.1 The two stages of the VERSF Born-rule programme
- 1.2 What this paper adds
- 1.3 The theory-architecture chain
- 1.4 Roadmap

2. The Five Routes as Projections of a Single Structure

- 2.1 The five derivations recapitulated
- 2.2 The unifying object: the admissibility framework
- 2.3 What each route enforces
- 2.4 Independence and unification: why both claims hold

3. The Admissibility Framework

- 3.1 The kernel A_0 – A_2 (inherited)
- 3.2 Derived structural conditions B_1 – B_2
- 3.3 Joint admissibility: the framework \mathfrak{A}
- 3.4 What \mathfrak{A} does and does not specify

4. The Admissible Probability Uniqueness Theorem

- 4.1 Statement
- 4.2 Logical architecture of the proof: substrate stage and rule stage
- *Substrate stage (Lemmas 1–3): deriving Hilbert structure from \mathfrak{A}*
- 4.3 Lemma 1 — Phase / $U(1)$ substrate from $A_1 \wedge A_2$
- 4.4 Lemma 2 — Local tomography from $A_0 \wedge A_1 \wedge B_1$ (selects \mathbb{C})
- 4.5 Lemma 3 — Compactness from $A_1 \wedge B_2$ (forces inner product)
- *Rule stage (Lemmas 4–7): deriving the Born rule on the substrate*
- 4.6 Lemma 4 — Magnitude-only dependence from $B_1 \wedge B_2$ (the bilinear core)
- 4.7 Lemma 5 — Trace-rule structure from $A_0 \wedge$ Route I
- 4.8 Lemma 6 — Exponent $p = 2$ from B_2 normalisation invariance
- 4.9 Lemma 7 — Iso-entropic stability from A_2
- 4.10 Closure: the intersection is a singleton
- 4.11 The killer line, made precise

5. The Five Routes as Single-Object Constraints

- 5.1 The unification table
- 5.2 What fails along each direction
- 5.3 Convergence becomes inevitability

6. Anticipated Referee Objections

- 6.1 The kernel-circularity objection
- 6.2 The hidden-postulate objection
- 6.3 The route-independence objection
- 6.4 The relation-to-Gleason objection
- 6.5 The contextuality-loop-hole objection
- 6.6 The "shared kernel \Rightarrow no independence" objection
- 6.7 The dimension-2 objection
- 6.8 The Lüders-rule objection

7. Relation to Prior Work

- 7.1 The earlier VERSF overdetermination paper
- 7.2 The kernel \rightarrow TPB companion paper
- 7.3 Standard reconstruction programmes

8. Falsification Conditions

9. Limitations and Scope

10. Conclusion

End matter

- Notation Summary
- Selected References

Notation Conventions

Symbol	Reading
\mathcal{A}	the set of admissible alternatives for a bounded subsystem
\mathcal{O}	the set of committed outcomes
$\mathcal{C} : \mathcal{A} \rightarrow \mathcal{O}$	the irreversible commitment map
\mathcal{G}	the group of admissible observer transformations
\mathcal{S}	a bounded subsystem

Symbol	Reading
\mathfrak{A}	the admissibility framework: $A0 \wedge A1 \wedge A2 \wedge B1 \wedge B2$
$\psi \in \mathcal{H}$	a pure state in the Hilbert space \mathcal{H}
$c_i = \langle i \psi \rangle$	
p_i	the probability assigned to outcome i
ρ	a density operator (mixed state)
E_i	a POVM effect for outcome i
$\{E_i\}$	a POVM, satisfying $\sum_i E_i = \mathbb{1}$ and $0 \leq E_i \leq \mathbb{1}$
ΔS_i	the entropic action associated with outcome i
λ	the inverse-temperature-like coefficient in the entropic action
\otimes	tensor product
\Rightarrow, \implies	implication
\wedge, \vee, \neg	logical conjunction, disjunction, negation

All amplitudes are complex-valued, $c_i \in \mathbb{C}$, with $\sum_i |c_i|^2 = 1$ in the standard normalisation.

1. Introduction: From Convergence to Necessity

1.1 The two stages of the VERSF Born-rule programme

The Born rule

$$p_i = |c_i|^2$$

is the most empirically successful and foundationally contested rule in quantum mechanics. Born postulated it. Gleason, Deutsch–Wallace, Zurek, Hardy, and Masanes–Müller derived it from various alternative postulates of contested status. Each successful reconstruction transferred the strangeness from the rule to the postulate; none dissolved it.

The VERSF programme has approached the rule in two stages.

Stage 1 — overdetermination. The companion paper *Why a Fact-Producing Universe Must Satisfy the Born Rule* established that five operationally independent derivations within VERSF converge exactly on the rule, sharing no proof step or premise beyond the kernel $A0 \wedge A1 \wedge A2$. This is a strong epistemic claim: the rule sits at the intersection of five logically independent constraints, any one of which would have to fail for the rule to be wrong. Convergence of this kind is not normal in foundational physics; it is, so far as we are aware, unique to the Born rule within the VERSF framework.

But overdetermination, even fivefold, leaves a question open. *Why* do five logically independent constraints converge on the same answer? Is the convergence a remarkable accident, or is it the visible signature of a deeper unity?

Stage 2 — inevitability. The present paper answers this question. The convergence is not an accident. The five routes are not five independent witnesses to the same fact; they are **five projections of a single underlying object** — the admissibility framework \mathfrak{A} — onto five different operational vocabularies. Each route tells you what \mathfrak{A} looks like from one direction. The Born rule is not merely the answer that all five routes happen to produce; it is the unique probability measure that \mathfrak{A} admits.

1.2 What this paper adds

We add four things.

First, we make the unifying object explicit. The admissibility framework $\mathfrak{A} = A_0 \wedge A_1 \wedge A_2 \wedge B_1 \wedge B_2$ is defined formally in §3. Its kernel $A_0 \wedge A_1 \wedge A_2$ is inherited from the broader VERSF foundational programme; the derived structural conditions B_1 (compositional admissibility) and B_2 (reversible pre-commitment dynamics) are not free postulates but follow from the kernel together with the requirement of a coherent compositional physical theory.

Second, we **derive Hilbert structure from \mathfrak{A}** , eliminating the most common shared input of single-route reconstructions. Three substrate lemmas (§§4.3–4.5) jointly establish that the admissible state space on any bounded subsystem is a finite-dimensional complex Hilbert space $\mathcal{H}(S)$: Lemma 1 gives the $U(1)$ phase substrate from $A_1 \wedge A_2$; Lemma 2 derives local tomography from $A_0 \wedge A_1 \wedge B_1$, selecting \mathbb{C} as the field over \mathbb{R} and \mathbb{H} ; Lemma 3 derives compactness of the reversible-dynamics group from $A_1 \wedge B_2$, forcing an invariant inner product. No external Hilbert-space presupposition remains.

Third, we prove the **Admissible Probability Uniqueness Theorem**: any probability assignment compatible with \mathfrak{A} must be $p_i = |\langle i|\psi\rangle|^2$ for pure states, with the natural extension to mixed states ρ via $p_i = \text{Tr}(E_i \rho)$. No alternative survives. The proof reuses existing mathematics: each of the five existing derivations is reinterpreted as a Lemma enforcing one structural constraint on \mathfrak{A} , and the joint satisfaction of the substrate lemmas (1–3) and rule lemmas (4–7) leaves the Born rule as the unique fixed point.

Fourth, we make the modal upgrade explicit. Convergence shows that the rule is *empirically robust* across multiple derivation paths. Inevitability shows that the rule is *structurally necessary*: more precisely, **the Born rule is the unique admissible probability measure within any fact-producing framework satisfying \mathfrak{A}** . Any alternative would have to violate at least one component of \mathfrak{A} , hence fail to be a fact-producing framework at all. The killer line at the close of §4 is:

Any deviation from quadratic probability either fails to be a fact-producing, composable, observer-invariant physical theory at all, or violates at least one admissibility constraint in \mathfrak{A} .

Each non-trivial deviation — from real-amplitude or quaternionic substrates, to non-compact symmetry groups, to non-bilinear kernels, to contextual rules, to non-quadratic exponents, to non-additive weightings, to phase-blind structures — fails at least one of the seven Lemmas. The Born rule is the unique rule that fails none.

1.3 The theory-architecture chain

A consequence of the substrate-derivation work in §4 is that the present result is no longer purely a Born-rule paper. It is a derivation of a layered theory architecture from a small set of fact-production principles. Schematically:

Facts ($A0 \wedge A1 \wedge A2$) ↓ kernel principles for observer-invariant content, finite admissibility, and irreversible commitment **Admissibility framework ($\mathfrak{A} = A0 \wedge A1 \wedge A2 \wedge B1 \wedge B2$)** ↓ kernel + derived conditions for compositional consistency and reversible pre-commitment dynamics **State structure (Hilbert space \mathcal{H} over \mathbb{C})** ↓ Lemmas 1–3: U(1) phase, local tomography, compact symmetry group with invariant inner product **Dynamics (unitary evolution)** ↓ B2 + Lemma 3: reversible dynamics act as a compact group of isometries **Probability (Born rule $p_i = |\langle i|\psi\rangle|^2$)** ↑ Lemmas 4–7: bilinear core, trace structure, exponent $p = 2$, iso-entropic stability

The chain runs strictly top-to-bottom: each layer is forced by the conjunction of layers above. No layer is a free postulate; each is a theorem from \mathfrak{A} . The downward direction is the deductive direction (kernel → kinematics → probability); the upward direction would be reconstruction (probability → kinematics → kernel), which is not what we do.

The strongest defensible claim of this paper is therefore not about the Born rule alone:

Quantum kinematics (complex Hilbert space, unitary dynamics) and the Born rule jointly emerge as the unique admissible structure consistent with $A0 \wedge A1 \wedge A2 \wedge B1 \wedge B2$.

Equivalently: any fact-producing, composable, observer-invariant physical theory whose pre-commitment dynamics preserve admissible distinctions has complex-Hilbert-space kinematics and Born-rule probability. The contestation, if any, must be at the layer of fact-production, finite admissibility, irreversible commitment, compositional consistency, or reversible pre-commitment dynamics — not at the kinematic or probabilistic layers, which are now derived rather than postulated.

1.4 Roadmap

§2 recapitulates the five routes and identifies the unifying object they each project. §3 specifies the admissibility framework formally. §4 proves the uniqueness theorem in seven lemmas (three substrate, four rule), with the substrate lemmas deriving Hilbert structure from \mathfrak{A} and the rule lemmas corresponding to the five Routes reinterpreted as single-object constraints. §5 makes the unification explicit through a tabulated correspondence between routes, constraints, and elimination conditions. §6 addresses eight referee objections, with particular attention to the kernel-circularity worry. §7 places the result in relation to prior work. §8 specifies in-principle falsification conditions. §9 states scope and limitations.

2. The Five Routes as Projections of a Single Structure

2.1 The five derivations recapitulated

Five derivations of the Born rule have been completed within the VERSF programme. We recapitulate them in compressed form; the full developments appear in the cited companion papers.

Route I — Closure / lattice structure. Starting from closure of admissible partitions on a measurement frame, with non-contextuality of frame functions imposed by observer invariance (A0), one obtains the Gleason–Busch result that any admissible probability assignment on the projector lattice must take the trace-rule form $\text{Tr}(E_i \rho)$. For pure states $\rho = |\psi\rangle\langle\psi|$ and projectors $E_i = |i\rangle\langle i|$, this reduces to $p_i = |\langle i|\psi\rangle|^2$.

Route II — Single-source consistency of records. Starting from the requirement that the committed record produced by a measurement be consistent across all admissible coarse-grainings — that the record assigned to a coarse outcome equals the sum of records assigned to its fine-grainings — one obtains the additivity condition that, combined with the inner-product structure on the underlying amplitude space, forces the quadratic measure.

Route III — Path correlations in discrete informational geometry. Starting from pairwise selection on path correlations under the discrete informational geometry of the BCB substrate, with normalisation under unitary evolution, one obtains the bilinear kernel structure. The quadratic form is the unique bilinear scalar invariant under the relevant symmetry group.

Route IV — Physical Admissibility Filter (PAF). Starting from the elimination of all alternative probability rules $p_i = |c_i|^p$ for $p \neq 2$ by direct application of admissibility constraints (gauge invariance, normalisation under composition, factorisation on product systems), one obtains $p = 2$ as the unique surviving exponent. Concrete counterexamples (Hadamard-type product-system constructions) demonstrate the failure of $p \neq 2$ cases.

Route V — Maximum-caliber entropic unfolding. Starting from the maximum-caliber inference principle applied to the entropic action functional $M_i = -\ln a_i + \lambda \Delta S_i$, with uniqueness forced by Cauchy's functional equation under composition, one obtains the generalised probability law

$$p_i \propto |c_i|^2 \cdot \exp(-\lambda \Delta S_i).$$

The Born rule is recovered in the iso-entropic limit $\Delta S_i = \text{const}$, which corresponds to standard measurement conditions.

Each derivation is a self-contained mathematical result, established in dedicated papers within the VERSF programme. Each starts from a logically distinct primitive — frame closure, record

consistency, path geometry, exponent admissibility, entropic inference — and arrives at the same probability measure. None invokes the Born rule as input.

2.2 The unifying object: the admissibility framework

The earlier overdetermination paper analysed the convergence of these five routes and identified their shared kernel — A0 (observer-invariant distinguishability), A1 (finite admissibility), A2 (irreversible commitment) — as the residual common dependency. The present paper goes one step further. We identify the *unifying object* that all five routes constrain.

That object is the **admissibility framework**, the conjunction

$$\mathfrak{A} := A0 \wedge A1 \wedge A2 \wedge B1 \wedge B2,$$

where the kernel A0–A2 is inherited from the broader VERSF programme and the derived structural conditions B1 (compositional admissibility) and B2 (reversible pre-commitment dynamics) follow from the kernel together with the requirement of having a coherent compositional physical theory. We motivate B1 and B2 in §3.

The five routes are then **five projections of \mathfrak{A} onto five different operational vocabularies**:

- Route I projects \mathfrak{A} onto **frame structure** — what valid measurement partitions look like.
- Route II projects \mathfrak{A} onto **record structure** — what valid commitment outputs look like.
- Route III projects \mathfrak{A} onto **path structure** — what valid amplitude correlations look like.
- Route IV projects \mathfrak{A} onto **exponent structure** — what valid scaling exponents look like.
- Route V projects \mathfrak{A} onto **inferential structure** — what valid entropic-inference assignments look like.

Each route makes one face of \mathfrak{A} operationally explicit and shows that, on that face, only the Born rule survives. The convergence of the five routes is then no longer an accident: it is the structural signature of the fact that they all face the same object.

2.3 What each route enforces

Within the unifying picture, each route enforces a specific admissibility constraint on the single object:

Route	Constraint enforced
I — Closure / lattice	Admissible partition structure (non-contextual frame functions)
II — Record consistency	Admissible weighting (additivity under coarse-graining)
III — Path correlations	Admissible selection (bilinear scalar invariance)
IV — Admissibility elimination	Admissible exponent ($p = 2$ forced)

Route	Constraint enforced
V — Entropic unfolding	Admissible measurement realisation (stability under commitment)

These are not five independent confirmations of the Born rule. They are five consistency conditions on a single object — admissible probability — and the Born rule is the unique measure satisfying all five conditions simultaneously.

2.4 Independence and unification: why both claims hold

A careful reader will already worry: if the routes are "five projections of one object", in what sense are they independent? We address this in detail at §6.3 and §6.6, but state the resolution here.

Independence and unification operate at different levels. The five routes are **operationally independent**: they share no proof step, no auxiliary lemma, no mathematical machinery beyond the kernel $A_0 \wedge A_1 \wedge A_2$. Any one of them can be developed and verified in isolation from the others. The five routes are also **structurally unified**: they constrain the same underlying object, the admissibility framework \mathfrak{A} . Each is a projection of \mathfrak{A} onto a distinct operational face.

Both claims are evidentially significant. Operational independence rules out shared-error explanations of the convergence: it cannot be the case that all five derivations are wrong because of a common bug in the proof machinery, because there is no common proof machinery. Structural unification explains why the convergence occurs in the first place: it is not a coincidence among independent witnesses but a necessity imposed by the shared object they constrain.

The analogy is with five independent measurements of the same physical quantity using different equipment and different methods. The independence is methodological (no shared apparatus or technique); the unification is ontic (they all measure the same quantity). Both are real, and both are evidentially significant.

3. The Admissibility Framework

We specify the admissibility framework formally. This section is unavoidably terse; it is the foundation for §4 and is best read alongside the kernel→TPB companion paper for full justification of the kernel principles.

3.1 The kernel A_0 – A_2 (inherited)

The admissibility framework rests on the VERSF kernel.

A0 — Observer-Invariant Distinguishability. Physical content consists only of distinctions invariant under the group \mathcal{G} of admissible observer transformations.

Rationale. If a "distinction" varies under the choice of observer transformation, no observer-independent fact corresponds to it; A0 simply makes precise what counts as physical content in any theory committed to observer-independence.

A1 — Finite Admissibility. Any bounded subsystem \mathcal{S} admits a finite set of admissible alternatives:

$$|\mathcal{A}(\mathcal{S})| < \infty.$$

Rationale. Bounded systems with infinite admissible alternatives generate divergences incompatible with the existence of stable physical facts; A1 is the minimal cutoff that makes commitment well-defined on bounded resources.

A2 — Irreversible Commitment. Physical facts are produced by an irreversible map $\mathcal{C} : \mathcal{A} \rightarrow \mathcal{O}$ satisfying:

- **many-to-one resolution** — the map collapses pre-commitment alternatives to committed outcomes;
- **non-invertibility** — the inverse \mathcal{C}^{-1} does not exist as an admissible operation;
- **positive entropy cost** — $\Delta S \geq k_B \ln 2$ per resolved bit (Landauer bound).

Rationale. A2 makes precise what it is for a fact to be produced. Without an irreversibility-producing step, "facts" cannot be distinguished from "potentialities" in any operationally meaningful way.

These three principles are inherited from the VERSF kernel and are taken here as given. The companion paper *Discrete Commitment and the Necessity of Tick–Bit Structure* derives the discrete tick–bit substrate from these principles; the present paper uses them as foundational input from which the admissibility framework is built.

3.2 Derived structural conditions B1–B2

Two derived conditions extend the kernel to a complete admissibility framework for probability. We emphasise: these are *derived* rather than postulated. Each follows from the kernel together with the requirement of having a coherent compositional physical theory.

B1 — Compositional Admissibility. For any two admissible subsystems \mathcal{S}_A and \mathcal{S}_B with joint system \mathcal{S}_{AB} , the probability assignment on \mathcal{S}_{AB} must satisfy a factorisation condition consistent with the kernel. Concretely, marginal probabilities on \mathcal{S}_A and \mathcal{S}_B obtained from \mathcal{S}_{AB} must equal the probability assignments computed directly on \mathcal{S}_A and \mathcal{S}_B :

$$\sum_b p(a, b) = p_A(a), \sum_a p(a, b) = p_B(b).$$

For independent subsystems with joint amplitude structure $c_{(a,b)} = c_a \cdot c_b$, the joint probability factorises:

$$p(a, b) = p_A(a) \cdot p_B(b).$$

Derivation from the kernel (operational form). We show that any violation of B1 produces an *operational* contradiction with A0, not merely a philosophical discomfort.

Suppose B1 fails: the probability assigned to subsystem \mathcal{S}_A obtained by marginalising over \mathcal{S}_B from the joint system \mathcal{S}_{AB} differs from the probability assigned to \mathcal{S}_A computed in isolation. Concretely, $\sum_b p_{AB}(a, b) \neq p_A(a)$ for some outcome a .

Consider two observers running identical measurement protocols on \mathcal{S}_A . Observer 1 treats \mathcal{S}_A as an isolated system and records outcome frequencies converging to $p_A(a)$. Observer 2 measures the joint system \mathcal{S}_{AB} and obtains the marginal frequency for a by tracing out \mathcal{S}_B ; her frequencies converge to $\sum_b p_{AB}(a, b)$. By hypothesis, these two limits differ.

But the outcomes recorded on \mathcal{S}_A are physical facts in the sense of A2 — they are committed records produced by \mathcal{C} . Their frequencies are therefore observer-invariant content under A0. The hypothesis $\sum_b p_{AB}(a, b) \neq p_A(a)$ makes outcome frequencies on \mathcal{S}_A depend on whether the observer chooses to consider \mathcal{S}_A in isolation or as embedded in \mathcal{S}_{AB} — a choice that is part of the observer's analytical decomposition, not part of any fact about \mathcal{S}_A itself. **This is observer-dependent outcome frequency, in direct violation of A0.**

B1 is therefore not an additional postulate but the unique condition under which subsystem outcome frequencies are observer-invariant. ■

B2 — Reversible Pre-Commitment Dynamics. Pre-commitment evolution — the dynamics on \mathcal{A} prior to action by \mathcal{C} — must preserve admissible distinctions. Operationally, the dynamics must be a representation of the symmetry group of admissible distinctions; for the kernel-compatible distinction structure of the standard quantum case, this is unitary:

$$U : \mathcal{A} \rightarrow \mathcal{A}, U^\dagger U = \mathbb{1}.$$

Derivation from the kernel (dichotomy form). Suppose pre-commitment dynamics \mathcal{U} fail to be distinguishability-preserving on the admissible alternatives \mathcal{A} . Then the cardinality of operationally distinguishable alternatives changes under \mathcal{U} . Two cases exhaust the failure modes.

Case 1 — Distinctions are destroyed ($|\mathcal{A}|$ decreases). Two alternatives that were operationally distinguishable before \mathcal{U} are mapped to a single alternative after \mathcal{U} . Operational distinguishability has been irreversibly lost; this *is* a commitment event by the definition of A2 (many-to-one resolution with positive entropy cost). But by hypothesis it occurred during pre-commitment evolution, not under the action of \mathcal{C} . Hence A2's role as the **unique** irreversibility-producing step is contradicted: commitment has leaked into the dynamics. We call this **premature commitment**.

Case 2 — Distinctions are created ($|\mathcal{A}|$ increases). A single alternative is mapped under \mathcal{U} to two previously indistinguishable alternatives. A new admissible distinction has appeared without any commitment event having occurred. This is **spontaneous fact generation**: a new distinction — a new admissible fact — has come into existence without action by \mathcal{C} . A2 is contradicted from the other direction, this time as the *necessary* condition for fact production (no fact arises except via \mathcal{C}). In addition, in any bounded subsystem, iterated application of distinction-creating dynamics violates the A1 finiteness bound.

Both failure modes contradict the kernel. Hence pre-commitment dynamics must be distinguishability-preserving:

$$|\mathcal{A}(\mathcal{U}\psi)| = |\mathcal{A}(\psi)| \text{ for all admissible } \psi.$$

For the kernel-compatible distinction structure (a complex inner-product space, justified in the kernel \rightarrow TPB and quantum-kinematics companion work), distinguishability preservation on the full state space is equivalent to unitarity. ■

B1 and B2 are not additional postulates introduced specifically to derive the Born rule. They follow from the kernel together with the elementary requirement of compositionality and the uniqueness of A2 as the irreversibility locus. We discuss the sense in which they are "derived rather than postulated" further at §6.2.

3.3 Joint admissibility: the framework \mathfrak{A}

The **admissibility framework** is the conjunction:

$$\mathfrak{A} := A0 \wedge A1 \wedge A2 \wedge B1 \wedge B2.$$

Any probability assignment compatible with \mathfrak{A} is called **admissible**. The central theorem of this paper (§4) is that the admissible probability measures form a single isomorphism class — namely, the Born rule.

3.4 What \mathfrak{A} does and does not specify

We are explicit about the content of \mathfrak{A} , since the kernel-circularity objection (§6.1) depends on it.

\mathfrak{A} specifies:

- what counts as observer-invariant content (A0);
- that bounded subsystems have finite admissible alternatives (A1);
- that facts are produced by irreversible commitment (A2);
- that probability assignments — *if they exist* — must factorise compositionally (B1);
- that pre-commitment dynamics must be reversible (B2).

\mathfrak{A} does not specify:

- the form of any probability rule;
- the existence of probability at all;
- the value of any probability;
- any quadratic, bilinear, or otherwise specific functional structure on amplitudes.

The Born rule is therefore not built into \mathfrak{A} in any direct sense. It emerges as the unique probability measure consistent with \mathfrak{A} once we ask the conditional question: *given that a probability assignment exists, what form must it take to be consistent with \mathfrak{A} ?*

The analogy with thermodynamics is exact. The laws of thermodynamics specify what counts as a thermodynamic state and what the structural relations between states are; they do not, in themselves, specify the form of the entropy function. But once the structural conditions are imposed, the entropy function is uniquely determined (up to additive and multiplicative constants). The Born rule has the same status here: it is the entropy-analogue of the admissibility framework, the unique measure consistent with the structure.

4. The Admissible Probability Uniqueness Theorem

4.1 Statement

Theorem (Admissible Probability Uniqueness). Let a physical theory satisfy the admissibility framework $\mathfrak{A} = A0 \wedge A1 \wedge A2 \wedge B1 \wedge B2$. Then any probability assignment over outcomes of an admissible measurement is uniquely determined by the state and effects, and takes the form

$$p_i = |\langle i|\psi\rangle|^2$$

for pure states $|\psi\rangle$, with the natural extension to mixed states ρ via $p_i = \text{Tr}(E_i \rho)$.

4.2 Logical architecture of the proof

The proof has the following architecture:

$$\mathfrak{A} \implies \{\text{Lemma 1, } \dots, \text{Lemma 7}\} \implies p_i = |\langle i|\psi\rangle|^2.$$

The seven lemmas split into two stages:

Stage A — substrate lemmas (Lemmas 1–3). Establish that the admissible state space is a finite-dimensional complex Hilbert space, derived entirely from \mathfrak{A} . This stage closes the gap that single-route reconstructions leave open: rather than presupposing Hilbert structure as input, we derive it.

- Lemma 1 (phase / $U(1)$ substrate from $A1 \wedge A2$) — amplitudes carry $U(1)$ phase.

- Lemma 2 (local tomography from $A0 \wedge A1 \wedge B1$) — composition is product-dimensional, selecting \mathbb{C} over \mathbb{R} and \mathbb{H} .
- Lemma 3 (compactness from $A1 \wedge B2$) — the symmetry group is compact, forcing an invariant inner product.

Stage B — probability-rule lemmas (Lemmas 4–7). Given the substrate, fix the probability rule.

- Lemma 4 (magnitude-only dependence from $B1 \wedge B2$) — $p_i = f(|c_i|^2)$.
- Lemma 5 (trace-rule from $A0 + \text{Route I}$, via Gleason–Busch on the substrate of Stage A) — $p_i = \text{Tr}(E_i \rho)$.
- Lemma 6 (exponent $p = 2$ from $B2$ normalisation, Route IV) — fixes $\alpha = 1$.
- Lemma 7 (iso-entropic stability from $A2$, Route V) — locates the Born rule as the iso-entropic limit of a deeper inference law.

The first implication is established lemma by lemma in §§4.3–4.9: each constraint in \mathfrak{A} forces one structural condition on either the substrate or the probability rule. The second implication is the **closure step** (§4.10): the joint satisfaction of all seven lemmas leaves the Born rule as the unique fixed point.

The seven lemmas reuse mathematics already established in the broader programme. Lemma 1 invokes the kernel→TPB substrate result; Lemma 2 uses dimension-counting arguments familiar from Hardy and Wootters; Lemma 3 uses Haar averaging on a compact group; Lemmas 4–7 are the existing five derivations restated as constraints on the substrate. The novelty is structural: the substrate lemmas eliminate Hilbert space as a presupposition, and the closure shows that the seven constraints jointly leave only one rule standing.

4.3 Lemma 1 — Phase / U(1) substrate from $A1 \wedge A2$

Lemma 1. $A1$ (finite admissibility) and $A2$ (irreversible commitment), through the discrete tick-bit substrate result of the companion kernel→TPB paper, imply that pre-commitment alternatives are organised into a discrete information-theoretic structure that admits a U(1) phase. The amplitudes carrying admissible content take values in \mathbb{C} rather than \mathbb{R} .

Proof sketch. The kernel→TPB result establishes that the kernel $A0 \wedge A1 \wedge A2$ forces a discrete tick-bit substrate at the per-worldline level. The structure of irreversible commitment events along worldlines, combined with finite admissibility per event, produces a continuous phase parameter on the pre-commitment alternatives — the holonomy of the discrete substrate. This phase is needed to maintain operational distinctness across reversible pre-commitment evolution while remaining consistent with finite per-event distinguishability. Without it, distinct pre-commitment trajectories that produce the same magnitudes could not be distinguished, contradicting $A0$. Hence the candidate amplitudes c_i must take values in a structure carrying U(1) phase — concretely, $c_i \in \mathbb{C}$.

What this rules out. Real-amplitude theories. Any candidate probability rule built on real-valued amplitudes alone fails to provide the holonomy structure required by $A1 \wedge A2$. ■

This Lemma provides the **U(1) substrate**: amplitudes carry a continuous phase parameter. Lemmas 2 and 3 extend this to the full complex Hilbert space structure.

4.4 Lemma 2 — Local tomography from A0 \wedge A1 \wedge B1 (selects C)

Lemma 2 (No-Alternative Theorem for Composition). Let \mathcal{S}_A and \mathcal{S}_B be admissible bounded subsystems with joint system \mathcal{S}_{AB} . Under A0, A1, and B1, the state space of \mathcal{S}_{AB} satisfies the **local tomography** condition: every state ρ_{AB} is uniquely determined by the joint statistics of local product measurements,

$$\rho_{AB,1} = \rho_{AB,2} \Leftrightarrow \text{Tr}(\rho_{AB,1} (E_i \otimes F_j)) = \text{Tr}(\rho_{AB,2} (E_i \otimes F_j)) \text{ for all admissible } E_i \in \mathcal{E}(\mathcal{S}_A), F_j \in \mathcal{E}(\mathcal{S}_B).$$

Any composition violating this condition falls into one of three failure modes — hidden global parameters, non-reconstructible states, or non-factorising marginals — each of which contradicts at least one component of \mathfrak{A} .

Proof (three-case exclusion). Set $N_A := |\mathcal{A}(\mathcal{S}_A)|$ and $N_B := |\mathcal{A}(\mathcal{S}_B)|$, both finite by A1. By the operational form of B1 (§3.2), the joint admissible-alternative set is the product:

$$\mathcal{A}(\mathcal{S}_{AB}) = \mathcal{A}(\mathcal{S}_A) \times \mathcal{A}(\mathcal{S}_B), |\mathcal{A}(\mathcal{S}_{AB})| = N_A \cdot N_B.$$

By A0, the observer-invariant content of \mathcal{S}_{AB} consists exactly of probability assignments over $\mathcal{A}(\mathcal{S}_{AB})$. The space of such assignments has dimension $N_A \cdot N_B - 1$ (subject to normalisation).

Local tomography is the condition that the state space of \mathcal{S}_{AB} has dimension exactly equal to $N_A \cdot N_B - 1$, with the additional condition that joint local statistics on product states factorise. Suppose for contradiction that local tomography fails. Then exactly one of the following three cases holds.

Case (C1) — Hidden global parameters. $\dim(\text{state space of } \mathcal{S}_{AB}) > N_A \cdot N_B - 1$.

Then there exist distinct states $\rho_1 \neq \rho_2$ on \mathcal{S}_{AB} giving identical probability assignments over $\mathcal{A}(\mathcal{S}_{AB})$ but differing in some additional parameter $\xi \in \Xi$. Since the admissible content of \mathcal{S}_{AB} is exhausted by $\mathcal{A}(\mathcal{S}_{AB})$ (by A1 + B1), ξ does not correspond to any admissible alternative.

Two sub-cases:

(C1a) ξ does not transform under \mathcal{G} . Then ξ is observer-dependent content present in the state but invisible to admissible measurement, contradicting A0's identification of physical content with observer-invariant distinguishability. ξ either represents non-physical surplus structure or, if it has any operational consequences, those consequences are observer-dependent — both ruled out.

(C1b) ξ transforms under \mathcal{G} but corresponds to no committed fact in \mathcal{O} . Then ξ is observer-invariant content not reducible to any committed record. By A2, all physical facts are produced by the commitment map $\mathcal{C} : \mathcal{A} \rightarrow \mathcal{O}$; observer-invariant content not produced by \mathcal{C} is content without a commitment route, which contradicts A2's status as the unique fact-producing step. ξ either fails to be physical content (ruled out by hypothesis that it is observer-invariant) or implicitly introduces a second commitment locus parallel to \mathcal{C} (ruled out by A2's uniqueness).

In both sub-cases, the state-space excess is incompatible with \mathfrak{A} . Hence $\dim(\text{state space}) \leq N_A \cdot N_B - 1$.

Case (C2) — Non-reconstructible states. $\dim(\text{state space of } \mathcal{S}_{AB}) < N_A \cdot N_B - 1$.

Then some normalised probability assignments $p : \mathcal{A}(\mathcal{S}_A) \times \mathcal{A}(\mathcal{S}_B) \rightarrow [0, 1]$ over the admissible-alternative set are *not* realised by any admissible state of \mathcal{S}_{AB} . Pick such an unrealised p with admissible marginals $p_A(a) := \sum_b p(a,b)$ and $p_B(b) := \sum_a p(a,b)$.

By A1 + B1, p_A and p_B are themselves admissible state assignments on \mathcal{S}_A and \mathcal{S}_B respectively. By the operational form of B1 (§3.2), any consistent assignment of marginal probabilities must arise from some admissible joint state — otherwise an observer measuring marginals on \mathcal{S}_A would record frequencies whose joint origin she cannot identify, while another observer measuring the same joint system records joint frequencies inconsistent with the first observer's marginals. Two observers therefore record outcome frequencies whose consistency cannot be reconciled by any admissible joint state. This is observer-dependent reconciliation of facts, in direct violation of A0 via the operational form of B1.

Hence every consistent probability assignment over $\mathcal{A}(\mathcal{S}_A) \times \mathcal{A}(\mathcal{S}_B)$ must be realised, giving $\dim(\text{state space}) \geq N_A \cdot N_B - 1$.

Case (C3) — Non-factorising marginals on product states. \dim equality holds, but joint local statistics fail to factorise on independent subsystem preparations.

Concretely: suppose ρ_A and ρ_B are admissible states of \mathcal{S}_A and \mathcal{S}_B respectively, each prepared independently. Let ρ_{AB} be the admissible joint state corresponding to " \mathcal{S}_A in state ρ_A , \mathcal{S}_B in state ρ_B , no correlation." B1 (factorisation form) requires:

$$\text{Tr}(\rho_{AB} (E_i \otimes F_j)) = \text{Tr}(\rho_A E_i) \cdot \text{Tr}(\rho_B F_j).$$

Suppose instead $\text{Tr}(\rho_{AB} (E_i \otimes F_j)) \neq \text{Tr}(\rho_A E_i) \cdot \text{Tr}(\rho_B F_j)$ for some admissible E_i , F_j . Then independent preparations produce correlated outcome frequencies — outcomes on \mathcal{S}_A statistically depend on outcomes on \mathcal{S}_B without any causal mechanism (by hypothesis the subsystems are uncorrelated). The correlation must therefore depend on the *embedding* of \mathcal{S}_A and \mathcal{S}_B into \mathcal{S}_{AB} rather than on their actual states.

But the embedding is observer-dependent: another observer treating \mathcal{S}_A in isolation records marginal frequencies $p_A(a)$ directly, with no embedding-induced correlations. Two observers therefore record different marginal frequencies on \mathcal{S}_A — one with embedding-induced correlations, one without. By the operational form of B1, this is observer-dependent outcome frequency, contradicting A0.

Equivalently: by A2, committed records on \mathcal{S}_A are facts; if their frequencies depend on the embedding rather than on ρ_A alone, the same physical setup yields different records for different observers, which is incompatible with A2's irreversible commitment producing observer-invariant outcomes. Hence factorisation must hold.

Closure of the three cases. The three cases (C1, C2, C3) are mutually exclusive and exhaustive: a tomographic failure means either too many state-space dimensions (C1), too few (C2), or the right number with anomalous correlations (C3). Each case violates at least one of A0, A1, A2, B1. Hence local tomography holds: $\dim(\text{state space of } \mathcal{S}_{AB}) = N_A \cdot N_B - 1$ with factorising marginals on independent products. ■

Corollary: \mathbb{C} as the unique amplitude field. Local tomography in finite dimension selects the amplitude field uniquely from the three associative normed division algebras ($\mathbb{R}, \mathbb{C}, \mathbb{H}$):

- *Real amplitudes* ($\mathbb{F} = \mathbb{R}$). For real Hilbert spaces, the symmetric tensor product satisfies $\dim_{\mathbb{R}}(\mathcal{H}_A \otimes_{\mathbb{R}} \mathcal{H}_B) > N_A \cdot N_B - 1$ due to the absence of phase: real density matrices on a composite system carry global structural parameters not capturable by local product POVMs (the antisymmetric components of correlations between non-commuting local observables). This is a Case (C1) failure.
- *Quaternionic amplitudes* ($\mathbb{F} = \mathbb{H}$). For quaternionic Hilbert spaces, the non-commutativity of \mathbb{H} produces $\dim_{\mathbb{H}}(\mathcal{H}_A \otimes_{\mathbb{H}} \mathcal{H}_B) < N_A \cdot N_B - 1$ because consistent quaternionic tensor products do not exist for general subsystems (Aaronson's argument: the ambiguity in left- vs right- multiplication structures eliminates compositional consistency). Some admissible marginal assignments have no joint-state realisation. This is a Case (C2) failure.
- *Complex amplitudes* ($\mathbb{F} = \mathbb{C}$). For complex Hilbert spaces, $\dim_{\mathbb{C}}(\mathcal{H}_A \otimes_{\mathbb{C}} \mathcal{H}_B) = (\dim_{\mathbb{C}} \mathcal{H}_A)(\dim_{\mathbb{C}} \mathcal{H}_B)$, and joint local statistics on product states factorise canonically. Cases (C1), (C2), (C3) are all avoided.

Only \mathbb{C} satisfies all three exclusions of Lemma 2 simultaneously. Hence the $U(1)$ phase of Lemma 1 extends to a full complex amplitude structure: $c_i \in \mathbb{C}$, and composite amplitude spaces are tensor products $\mathcal{H}_A \otimes_{\mathbb{C}} \mathcal{H}_B$ over the complex field.

What this rules out. Real-amplitude theories, quaternionic-amplitude theories, and the broader class of generalised probabilistic theories (GPTs) that satisfy A0–A2 and B1 in some form but

allow non-tomographic composition. The amplitude field is \mathbb{C} , derived from \mathfrak{A} by exclusion of all alternatives, not assumed. ■

This Lemma closes the **field-selection gap**: \mathbb{C} is forced via three independent exclusions, each tied to a specific component of \mathfrak{A} .

4.5 Lemma 3 — Compactness from A1 \wedge B2 (forces inner product)

Lemma 3. Under A1 (finite admissibility) and B2 (reversible pre-commitment dynamics), the group K of admissible reversible transformations on any bounded subsystem \mathcal{S} is compact. Consequently, the admissible state space carries a K -invariant inner product, and admissible reversible dynamics act unitarily.

Proof sketch. A1 specifies $|\mathcal{A}(\mathcal{S})| = N < \infty$. The space of admissible states on \mathcal{S} — probability assignments over $\mathcal{A}(\mathcal{S})$ — is the $(N - 1)$ -simplex in \mathbb{R}^N , which is closed and bounded. In any topology consistent with the kernel (e.g., total-variation distance on the simplex, or any norm-induced topology on the enclosing real vector space), this state space is **compact**.

By B2 (dichotomy form, §3.2), K acts on the admissible state space by distinguishability-preserving bijections. K is therefore a subgroup of the homeomorphism group of the admissible state space.

Suppose for contradiction that K is non-compact. Then K contains an unbounded one-parameter subgroup $K_t = \exp(tX)$, $t \in \mathbb{R}$, generated by some admissible generator X . Acting on a fixed admissible state ψ produces an orbit:

$$O(\psi) = \{K_t \psi : t \in \mathbb{R}\}.$$

By the unboundedness of K_t in any operational metric, this orbit is unbounded in the state-space metric. We now make the bridge to A1 explicit.

The bridge: unbounded distinguishability growth. The chain has three links — geometry \rightarrow distinguishability \rightarrow admissibility — and we make each one explicit. **Unbounded group action implies arbitrarily fine distinguishability resolution, contradicting the finite capacity bound of A1.** In detail: an unbounded orbit in the operational metric means that the operational distinguishability between $K_t \psi$ and $K_s \psi$ is unbounded as $|t - s| \rightarrow \infty$. By B2, distinguishability is preserved along the orbit, so each $K_t \psi$ is operationally distinguishable from $K_s \psi$ whenever $t \neq s$ (otherwise the dynamics would have collapsed two formerly-distinct states, violating B2). Hence the orbit $O(\psi)$ carries a continuum of pairwise operationally distinguishable admissible states.

But A1 bounds the number of operationally distinguishable admissible states in any bounded subsystem: $|\mathcal{A}(\mathcal{S})| = N < \infty$. **Unbounded distinguishability growth in a bounded subsystem contradicts A1's finite-distinguishability bound.** Specifically, after sufficiently many iterations along K_t , the orbit produces more than N pairwise-distinguishable admissible states, which cannot all be elements of $\mathcal{A}(\mathcal{S})$.

Equivalently in the dual formulation: the orbit either (i) exits the admissible state space — contradicting B2 directly — or (ii) packs a continuum of operationally distinct admissible states into a finite set, contradicting A1.

Both consequences contradict the kernel. Hence K is compact. ■

Corollary: invariant inner-product structure. A compact group K acting linearly on a finite-dimensional complex vector space (the field is \mathbb{C} by Lemma 2) admits a K -invariant Hermitian inner product, obtained by Haar averaging:

$$\langle v, w \rangle_K := \int_K \langle k \cdot v, k \cdot w \rangle d\mu(k),$$

where μ is the (normalised) Haar measure on K and $\langle \cdot, \cdot \rangle$ is any reference inner product. K acts by isometries of $\langle \cdot, \cdot \rangle_K$, hence is a subgroup of the unitary group $U(N)$.

Combined consequence (Lemmas 1 + 2 + 3). The admissible state space on any bounded subsystem \mathcal{S} is a finite-dimensional complex inner-product space — that is, a finite-dimensional **Hilbert space $\mathcal{H}(\mathcal{S})$** — on which the admissible reversible dynamics act unitarily. *Hilbert structure is therefore not inherited from external work but derived from \mathcal{U} .* This closes the gap that single-route reconstructions leave open. ■

4.6 Lemma 4 — Magnitude-only dependence from $B1 \wedge B2$ (the bilinear core)

Lemma 4. $B1$ (compositional admissibility) combined with $B2$ (reversible pre-commitment dynamics), acting on the Hilbert structure derived in Lemmas 1–3, forces the probability assignment to depend only on the magnitudes $|c_i|$ of the amplitudes, not on their phases. Specifically, any admissible probability rule has the form

$$p_i = f(|c_i|^2)$$

for some scalar function $f : [0, 1] \rightarrow [0, 1]$.

Proof sketch. Two constraints combine.

Phase invariance. $B2$ requires pre-commitment dynamics to be unitary on the Hilbert structure of Lemma 3. The global phase transformation $|\psi\rangle \mapsto e^{i\theta} |\psi\rangle$ is unitary, hence admissible. Under this transformation $c_i \mapsto e^{i\theta} c_i$ for all i , leaving $|c_i|^2$ invariant. Probabilities must be invariant under admissible unitary transformations (by $B2$), so probabilities cannot depend on the global phase. By a basis-rotation argument, the local phase of each c_i — i.e. the relative phase to a fixed reference — also cannot enter p_i in any way that survives composition with arbitrary unitaries on a tensor factor. Hence p_i depends only on $|c_i|^2$ (and possibly on the magnitudes of other components, addressed by Lemma 5 below).

Compositional admissibility. Consider two independent subsystems with amplitudes c_a and c_b , joint amplitude $c_{(a,b)} = c_a \cdot c_b$ on the tensor product $\mathcal{H}_A \otimes \mathcal{H}_B$ (well-defined by Lemma 2). $B1$ requires factorisation:

$$p_{(a,b)} = p_a \cdot p_b.$$

If $p_i = f(|c_i|^2)$, then $f(|c_a|^2 \cdot |c_b|^2) = f(|c_a|^2) \cdot f(|c_b|^2)$. This is the multiplicative Cauchy functional equation:

$$f(xy) = f(x) \cdot f(y), \quad x, y \in [0, 1].$$

The continuous solutions are $f(x) = x^\alpha$ for some $\alpha \geq 0$. The trivial case $\alpha = 0$ gives constant probability and is excluded by normalisation. Hence

$$p_i = |c_i|^{2\alpha}$$

for some $\alpha > 0$. Lemma 6 will fix $\alpha = 1$.

What this rules out. Phase-dependent probability rules. Higher-order kernels of the form $K(c, c^*, c, c^*, \dots)$ with explicit phase structure that survives unitary invariance. Non-multiplicative factorisation laws on product systems. ■

This Lemma provides the **bilinear core**: the probability is a power of $|c_i|^2$, not of c_i directly or of any structure involving phase. It does not yet specify the exponent.

4.7 Lemma 5 — Trace-rule structure from A0 \wedge Route I

Lemma 5. A0 (observer invariance) combined with the closure of admissible partitions (Route I in lattice form) forces the probability assignment to be **non-contextual**: $p(i | \rho)$ depends only on ρ and E_i , not on the other elements of the POVM or the physical implementation of the measurement. By Gleason–Busch, in dimension $d \geq 3$, any non-contextual probability assignment on the lattice of admissible projectors (or, with the Busch extension, on the space of admissible effects) must take the trace-rule form

$$p(E | \rho) = \text{Tr}(E \rho).$$

Status of the Hilbert structure. Lemma 5 invokes Gleason–Busch on a Hilbert lattice. **Hilbert structure is now derived rather than inherited**: Lemmas 1, 2, and 3 jointly establish that the admissible state space is a finite-dimensional complex Hilbert space $\mathcal{H}(\mathcal{S})$, with reversible dynamics acting unitarily. The chain is:

A0 \wedge A1 \wedge A2 \rightarrow Lemma 1 (U(1) phase, \mathbb{C} -amplitudes) A0 \wedge A1 \wedge B1 \rightarrow Lemma 2 (local tomography, \mathbb{C} as field, tensor products) A1 \wedge B2 \rightarrow Lemma 3 (compactness, invariant inner product)

$\implies \mathcal{H}(\mathcal{S})$ is a finite-dimensional complex Hilbert space \implies Lemma 5 invokes Gleason–Busch on this derived structure.

No external presupposition of Hilbert space remains in the proof. The framework is closed: every input to Gleason–Busch is itself a consequence of \mathfrak{A} .

Proof sketch. A0 specifies that physical content consists only of distinctions invariant under \mathcal{G} . Suppose, for contradiction, that the probability $p(i | \rho)$ depended on the choice of POVM beyond its dependence on the specific effect E_i . The POVM choice — as opposed to the committed outcome and the effect — is part of the measurement context; it does not by itself carry observer-invariant content (the choice of measurement apparatus is observer-dependent). Hence dependence of p_i on the POVM-as-context smuggles non-content-bearing structure into the physical content, in violation of A0.

Therefore $p(i | \rho)$ is non-contextual. By Gleason's theorem (extended by Busch and by Caves, Fuchs, Manne, Renes to POVMs), in Hilbert dimension $d \geq 3$, the unique non-contextual probability assignment on the lattice of effects compatible with normalisation is the trace pairing $p(E | \rho) = \text{Tr}(E \rho)$.

What this rules out. Contextual probability assignments that depend on the physical implementation of the measurement, on the unmeasured elements of a POVM, or on the history of how the measurement was constructed. We address $d = 2$ separately at §6.7. ■

This Lemma provides the **measure structure**. Combined with the bilinear core of Lemma 4, the trace pairing for pure states $\rho = |\psi\rangle\langle\psi|$ and projectors $E_i = |i\rangle\langle i|$ reduces to

$$p_i = |\langle i|\psi\rangle|^{2\alpha},$$

with α from Lemma 4 still undetermined.

4.8 Lemma 6 — Exponent $p = 2$ from B2 normalisation invariance

Lemma 6. Among candidate forms $p_i = |c_i|^p$, only $p = 2$ is preserved under arbitrary unitary evolution required by B2 (acting on the Hilbert structure of Lemmas 1–3). Hence the exponent α in Lemma 4 is fixed: $\alpha = 1$, equivalently $p = 2$.

Proof sketch. This is the Physical Admissibility Filter (Route IV). The argument is constructive. Consider a Hadamard-type unitary acting on a 2×2 product system that mixes amplitudes:

$$H = (1/\sqrt{2}) \cdot ((1, 1), (1, -1)).$$

Under H , normalised input amplitudes (c_1, c_2) with $\sum |c_i|^2 = 1$ transform to (c_1', c_2') with $\sum |c_i'|^2 = 1$ (unitarity preserves the 2-norm). For the candidate rule $p_i = |c_i|^p$, normalisation requires $\sum |c_i|^p = 1$ both before and after.

Direct computation shows that $\sum |c_i|^p$ is preserved under arbitrary unitaries if and only if $p = 2$. Concretely, for the input $(c_1, c_2) = (1, 0)$:

$$\text{before: } \sum |c_i|^p = 1^p + 0^p = 1. \text{ after: } c_i' = (1/\sqrt{2}, 1/\sqrt{2}), \sum |c_i'|^p = 2 \cdot (1/\sqrt{2})^p = 2^{1-p/2}.$$

This equals 1 if and only if $p = 2$. For any $p \neq 2$, normalisation drifts under the Hadamard unitary, contradicting B2.

Extension to product systems and arbitrary unitaries proceeds by the same argument: any non-trivial mixing unitary fails normalisation conservation for $p \neq 2$.

What this rules out. All exponents $p \neq 2$. The only exponent compatible with the unitary invariance forced by B2 is $p = 2$. Combined with Lemmas 4 and 5:

$$p_i = |\langle i|\psi\rangle|^2 \text{ (for pure states), } p(E|\rho) = \text{Tr}(E\rho) \text{ (for general states and effects). } \blacksquare$$

4.9 Lemma 7 — Iso-entropic stability from A2

Lemma 7. Under A2 (irreversible commitment) with the entropic action functional $M_i = -\ln a_i + \lambda \Delta S_i$, the maximum-caliber inference principle yields the generalised probability law

$$p_i \propto |c_i|^2 \cdot \exp(-\lambda \Delta S_i),$$

which reduces to the Born rule $p_i = |c_i|^2 / Z$ in the iso-entropic limit $\Delta S_i = \text{const}$.

Proof sketch. This is Route V. Maximum-caliber inference applied to the entropic action functional, with uniqueness imposed by Cauchy's functional equation under composition, produces the displayed law. In standard quantum measurements the engineered conditions ensure ΔS_i is approximately constant across outcomes (the measurement apparatus is designed so that the entropic cost of resolving each outcome is the same), giving the Born rule as the iso-entropic limit.

What this rules out. Probability rules that are quadratic but unstable under the entropic structure of measurement processes — i.e. rules of the form $p_i = |c_i|^2$ that fail to reduce from a deeper inference principle in the appropriate limit. \blacksquare

This Lemma provides **stability under commitment**: the Born rule is robust as the iso-entropic limit of the deeper inference law. The non-iso-entropic regime predicts deviations from the Born rule of the form $p_i \propto |c_i|^2 \cdot \exp(-\lambda \Delta S_i)$, which would represent a *strengthening* rather than a falsification of the framework (see §8).

4.10 Closure: the intersection is a singleton

We now assemble the proof. Lemmas 1–7 establish:

Lemma	Forces	From	Stage
1	\mathbb{C} -amplitudes carrying U(1) phase	A1 \wedge A2	Substrate
2	Local tomography; field is \mathbb{C} (not \mathbb{R} , not \mathbb{H})	A0 \wedge A1 \wedge B1	Substrate
3	Compact symmetry group; invariant inner product	A1 \wedge B2	Substrate
4	$p_i = f(c_i ^2)$, with $f(x) = x^\alpha$ a power	B1 \wedge B2	Rule
5	Trace-rule measure structure	A0 \wedge Route I	Rule

Lemma	Forces	From	Stage
6	Exponent $\alpha = 1$ (equivalently $p = 2$)	B2 normalisation, Route IV	Rule
7	Iso-entropic stability	$A2 \wedge$ Route V	Rule

Stage A (Lemmas 1–3) establishes that the admissible state space is a finite-dimensional complex Hilbert space $\mathcal{H}(\mathcal{S})$ on which B2 acts unitarily. Stage B (Lemmas 4–7) acts on this derived substrate and fixes the probability rule.

The conjunction leaves exactly one surviving form for pure states:

$$p_i = |\langle i|\psi\rangle|^2,$$

with the natural extension to general states and effects via the trace pairing $p_i = \text{Tr}(E_i \rho)$ from Lemma 5.

No alternative satisfies all seven Lemmas simultaneously. The intersection is a singleton. ■

4.11 The killer line, made precise

Killer line. Any deviation from quadratic probability either fails to be a fact-producing, composable, observer-invariant physical theory at all, or violates at least one admissibility constraint in \mathfrak{A} .

We make this precise. A probability rule that is:

- **not phase-coherent** (no $U(1)$ substrate, real-valued amplitudes only) \rightarrow violates $A1 \wedge A2$ via Lemma 1;
- **not locally tomographic** (composite states carry hidden global parameters, or are non-reconstructible from local statistics, or have non-factorising marginals on product preparations) \rightarrow violates $A0 \wedge A1 \wedge B1$ via Lemma 2 (one of three exclusion cases C1, C2, C3); equivalently, the underlying field is \mathbb{R} or \mathbb{H} rather than \mathbb{C} , or the theory is a generalised probabilistic theory beyond standard quantum mechanics;
- **not compactly symmetric** (admissible reversible dynamics form a non-compact group) \rightarrow violates $A1$ via Lemma 3; equivalently, the state space lacks an invariant inner product;
- **phase-dependent** (depends on global or relative phase in non-invariant ways) \rightarrow violates $B2$ via Lemma 4;
- **non-multiplicative on product systems** (fails $B1$ factorisation) \rightarrow violates $B1$ via Lemma 4;
- **contextual** (depends on POVM choice beyond the specific effect) \rightarrow violates $A0$ via Lemma 5;
- **non-bilinear** (involves higher-order kernels not reducible to $f(|c|^2)$) \rightarrow violates $B1 \wedge B2$ via Lemmas 4–5;
- **of exponent $p \neq 2$** \rightarrow violates $B2$ normalisation via Lemma 6;

- **unstable under iso-entropic commitment** (does not reduce from a deeper inference principle in the appropriate limit) → violates A2 via Lemma 7.

The Born rule is the unique probability assignment compatible with \mathfrak{A} . Every alternative probability rule fails at least one constraint, by at least one Lemma. ■

5. The Five Routes as Single-Object Constraints

The proof in §4 changes how we should read the five derivations. They are not five independent confirmations of the Born rule; they are five constraints on the admissibility framework, each enforcing one of the four rule-stage Lemmas (Lemmas 4–7) of the uniqueness theorem. The three substrate lemmas (Lemmas 1–3) operate upstream of the Routes, deriving the Hilbert structure on which the Routes act.

5.1 The unification table

The seven lemmas split into **substrate lemmas** (Lemmas 1–3) which derive the Hilbert structure, and **rule lemmas** (Lemmas 4–7) which fix the probability assignment on that structure. The five Routes correspond to the rule lemmas:

Route	Enforces (Lemma)	Step in proof	What fails without it
I — Closure	Non-contextual frame functions (Lemma 5)	Trace-rule measure	Contextual rules survive
II — Record consistency	Additivity under coarse-graining (Lemmas 4–5)	Magnitude-only dependence + trace structure	Non-additive assignments survive
III — Path correlations	Bilinear scalar invariance (Lemma 4)	Magnitude-only dependence	Higher-order kernels survive
IV — Admissibility elimination	Exponent $p = 2$ (Lemma 6)	Power fixed	$p \neq 2$ forms survive
V — Entropic unfolding	Iso-entropic stability (Lemma 7)	Stability under commitment	Born rule survives only as restricted special case

The substrate lemmas (1–3) are upstream of all five Routes: every Route presupposes some structure on the amplitude space, and the substrate lemmas derive that structure from \mathfrak{A} . Without Lemmas 1–3, the Routes would each need to import Hilbert structure as an external assumption (as standard reconstructions do); with Lemmas 1–3, they act on a substrate that is itself a consequence of \mathfrak{A} .

5.2 What fails along each direction

Each route, taken alone, leaves room for alternative probability structures along the directions the other routes constrain. A rule satisfying Route I alone (non-contextual trace pairing) is consistent with non-quadratic exponents; a rule satisfying Route IV alone ($p = 2$) is consistent with contextual measure structures; a rule satisfying Route V alone (entropic stability) is consistent with non-bilinear kernels. Only the *joint* satisfaction — joint admissibility — eliminates all alternatives.

This is the precise sense in which a single-route reconstruction (Gleason, Deutsch–Wallace, Zurek, Hardy, Masanes–Müller) is structurally weaker than the joint admissibility reconstruction. A single route forces the rule along its own direction but leaves open whether non-Born rules survive along the other directions; only the joint satisfaction closes all directions simultaneously.

5.3 Convergence becomes inevitability

The earlier overdetermination paper showed that the Born rule sits at the intersection of five logically independent constraints. The present paper shows that this intersection is **unique**: no alternative probability rule lies anywhere in the admissibility framework. The Born rule is not one rule among several admissible options; it is the only admissible rule.

This is the modal upgrade. **Convergence** says: five routes happen to give the same answer. **Inevitability**, made precise, says:

The Born rule is the unique admissible probability measure within any fact-producing framework satisfying \mathfrak{A} .

We use the term "inevitable" throughout the paper as shorthand for this technical claim. It is not a claim that the Born rule holds in all logically possible universes, nor that \mathfrak{A} is the only possible framework. It is the claim that *within* \mathfrak{A} — within any framework committed to observer-invariance, finite admissibility, irreversible commitment, compositional consistency, and reversible pre-commitment dynamics — the Born rule is the unique admissible measure. Any other rule would violate at least one constraint and therefore fail to be admissible *as a probability rule for a fact-producing physics* at all.

6. Anticipated Referee Objections

A unification result of this kind invites scrutiny on multiple fronts. We anticipate the most pointed objections and respond directly.

6.1 The kernel-circularity objection

You've just hidden the Born rule in A0–A2 and B1–B2. The framework is constructed to force the Born rule, then surprised to discover that it does.

This is the most dangerous objection. We address it directly.

The kernel does not specify probability. A0 specifies what counts as observer-invariant content. A1 specifies that bounded subsystems have finite admissible alternatives. A2 specifies that facts are produced by irreversible commitment. None of these mentions probability, the form of any probability rule, or the existence of probability at all. A0 is about *what is physical*; A1 is about *how much is admissible*; A2 is about *what it is to produce a fact*. None is about *how to weight outcomes*.

The derived conditions do not specify probability either. B1 specifies that probability assignments — *if they exist* — must factorise compositionally. B2 specifies that pre-commitment dynamics must be reversible. Both are conditional on probability already being present in the theory; neither defines what the probability rule is, and neither presupposes any specific functional form.

Probability emerges as the unique measure consistent with the structure. The admissibility framework specifies what counts as a fact ($A0 \wedge A2$), what persists ($A1 \wedge A2$), how systems compose (B1), and how they evolve before commitment (B2). The theorem then asks: *given that a probability assignment exists at all, what form must it take to be consistent with this structure?* The answer — $p_i = |c_i|^2$ — is forced, not built in.

Test: drop probability from the framework. A reader can verify the non-circularity directly. Drop the conditional assumption that a probability rule exists. The framework still says everything it said before: that physical content is observer-invariant, that bounded systems have finite admissible alternatives, that facts are produced by irreversible commitment, that compositional structure factorises (when it exists), and that pre-commitment dynamics is reversible. The framework is a complete specification of admissible structure that mentions probability nowhere.

The analogy is with thermodynamics. The laws of thermodynamics specify what counts as a thermodynamic state and what the structural relations between states are; they do not, in themselves, specify the form of the entropy function. But once the structural conditions are imposed, the entropy function is uniquely determined (up to additive and multiplicative constants). The Born rule has the same status here: it is the entropy-analogue of the admissibility framework, the unique measure consistent with the structure. The fact that thermodynamic structure forces a unique entropy function is not evidence that entropy is "smuggled into" the laws of thermodynamics; it is evidence that the laws have content.

6.2 The hidden-postulate objection

B1 and B2 look like postulates introduced specifically to get the Born rule. Without them, the kernel A0–A2 alone is too weak.

B1 and B2 are not postulates introduced for the present derivation. The full derivations are given in §3.2 (operational form for B1, dichotomy form for B2); we summarise the structure here.

B1 is forced by A0 in compositional contexts. Any violation of B1 produces an observable inconsistency: two observers running identical measurement protocols on the same subsystem \mathcal{S}_A — one in isolation, the other as a marginalisation from a joint system — record different outcome frequencies. Since outcome frequencies on \mathcal{S}_A are committed records (facts in the sense of A2), they must be observer-invariant under A0. Hence B1 is forced. This is an *operational* contradiction, not a philosophical preference.

B2 is forced by the uniqueness role of A2. Any violation of distinguishability preservation by pre-commitment dynamics falls into one of two cases: distinctions are destroyed (premature commitment, contradicting A2's role as the *unique* irreversibility-producing step), or distinctions are created (spontaneous fact generation without action by \mathcal{C} , contradicting A2's role as the *necessary* condition for fact production; also violating A1's finiteness bound under iteration). Both cases contradict the kernel. Hence B2 is forced.

The challenge: name a different B1, B2. The cleanest test of the no-free-postulate claim is constructive. If B1 and B2 were free postulates, alternative versions of them would be consistent with A0–A2 and would lead to alternative probability rules. We invite challengers to specify such alternatives. Concretely: relaxing B1 produces theories with observer-dependent outcome frequencies (operational violation of A0); relaxing B2 produces theories with multiple irreversibility loci or spontaneous fact generation (contradiction with A2). Neither relaxation produces a different admissible probability rule; both produce theories that fail to be coherent fact-producing frameworks at all.

6.3 The route-independence objection

You claim the five routes are "logically independent" but reinterpret them as constraints on a "single object". If they're projections of the same object, in what sense are they independent?

The independence claim and the projection claim refer to different levels.

At the **operational** level, the five routes are independent: they share no proof step, no premise, and no mathematical machinery beyond the kernel $A0 \wedge A1 \wedge A2$. Route I uses lattice theory and Gleason–Busch; Route II uses record-consistency arguments; Route III uses path-integral structure; Route IV uses normalisation analysis under specific unitaries; Route V uses maximum-caliber inference. Any one of them can be developed and verified without invoking the others.

At the **structural** level, the five routes constrain a single underlying object (the admissibility framework). They are projections in the precise sense that each makes one face of \mathfrak{U} operationally explicit.

The two readings are compatible. Compare: five independent experiments measuring different properties of the same physical object are *experimentally independent* (different equipment, different methods) but *constrain the same object* (the object's properties). The independence is methodological; the unification is structural. Both are real, and both are evidentially significant. Methodological independence rules out shared-error explanations of the convergence; structural unification explains why the convergence occurs in the first place.

6.4 The relation-to-Gleason objection

Route I just is Gleason's theorem. The whole derivation is then Gleason plus auxiliary scaffolding. What's new?

Route I is Gleason–Busch, restated within the VERSF admissibility framework. We do not claim to reprove Gleason's theorem. What is new is two-fold.

First, the present result embeds Gleason in a wider unification — and now also derives its substrate. Gleason alone shows that, *given* a Hilbert lattice with non-contextual frame functions, the trace rule is forced. The present paper goes further: Lemmas 1–3 (§§4.3–4.5) **derive** the Hilbert lattice structure itself from \mathfrak{A} , using local tomography to fix the field as \mathbb{C} and compactness to fix the inner product. Lemma 5 then invokes Gleason–Busch on this *derived* substrate, not on a presupposed one. Gleason answers "given Hilbert space, what's the rule?"; the present unification answers "why Hilbert space and why this rule?" — and answers both from \mathfrak{A} alone.

Second, the present result closes loopholes Gleason leaves open. Gleason's theorem requires $d \geq 3$ and is silent on POVMs without the Busch extension; it does not rule out $p \neq 2$ alternatives that violate non-contextuality (these are excluded only by additional assumptions); it does not address compositional admissibility on product systems; and it presupposes a complex Hilbert space rather than deriving the field. The present framework closes each of these loopholes from independent directions: Lemma 6 (Route IV) closes the $p \neq 2$ loophole via normalisation under unitary mixing; B1 plus Lemma 2 close the compositional loophole and select \mathbb{C} over \mathbb{R} and \mathbb{H} ; Lemma 3 forces the inner product via compactness; the $d = 2$ case is addressed by Lemmas 4 and 6 which do not require $d \geq 3$ (§6.7). The unified result is therefore strictly stronger than Gleason alone.

6.5 The contextuality-loophole objection

Could there be a contextual probability rule — one that depends on the full POVM context — that satisfies the other four constraints?

This is the cleanest version of the contextuality concern. The answer is no, for two reasons.

First, A0 (observer invariance) directly forbids contextual probability assignments at the level of admissible content. A probability that depends on the choice of POVM, beyond its dependence on the specific effect E_i , requires the POVM choice to carry observer-invariant content beyond what the effect itself carries. But by $A0 \wedge A2$, observer-invariant content is exhausted by the committed facts; the POVM choice is part of the measurement context (an observer-dependent feature of how the apparatus is set up), not part of the committed content. Hence contextual probability assignments smuggle non-content-bearing structure into the physical content, in violation of A0.

Second, B1 (compositional admissibility) provides a backup exclusion. Even if A0 were weakened, contextual probability rules generically fail to factorise on product systems because

the context-dependence does not respect tensor-product structure. The Hadamard counterexamples used in Route IV's exclusion of $p \neq 2$ are easily extended to exclude contextual rules by the same compositional argument: a contextual rule on \mathcal{S}_A and \mathcal{S}_B will generally fail to give consistent marginals on \mathcal{S}_A from the joint system \mathcal{S}_{AB} .

The contextuality loophole is therefore closed twice over — by A0 directly, and by B1 indirectly through compositional failure.

6.6 The "shared kernel \Rightarrow no independence" objection

You admit all five routes share the kernel A0–A2. Doesn't that destroy the independence claim?

The kernel is shared input; the derivation paths are independent. This is no more an objection to independence than the fact that all five experiments use the same physical apparatus (the universe) is an objection to their methodological independence.

More precisely: independence in the sense relevant to overdetermination is the absence of *shared derivation steps*, not the absence of *shared inputs*. Shared derivation steps would create the possibility that a single bug propagates across all five routes (the shared-error worry). Shared inputs do not: a single error in the kernel would falsify all five derivations equally, but that is exactly the *correct* behaviour — if the kernel is wrong, the framework's claim about probability is wrong. The independence claim addresses a more specific worry, namely that the routes might appear to converge only because they share intermediate machinery; this worry is dispatched by their derivational independence.

Furthermore, the kernel itself is independently motivated. A0 follows from the requirement of observer-invariance in any physical theory; A1 follows from the requirement of bounded resources; A2 follows from the existence of physical facts as distinct from potentialities. The kernel is therefore not a free parameter chosen to make the routes converge; it is the minimal specification of what it is to have an observer-invariant, finite-resource, fact-producing theory.

6.7 The dimension-2 objection

Gleason–Busch fails in $d = 2$. Doesn't that leave a loophole for non-Born rules in qubit systems?

Lemma 5 (the trace-rule lemma) invokes Gleason–Busch and therefore requires $d \geq 3$. The $d = 2$ case is addressed by the other Lemmas, which do not depend on Gleason–Busch.

Specifically: in $d = 2$, the substrate lemmas (Lemmas 1–3) still derive complex Hilbert structure on the qubit (the local-tomography argument of Lemma 2 applies to qubits as subsystems of two-qubit systems where $N_A \cdot N_B \geq 4 \geq 3$, so the dimension restriction does not bite). Lemma 4 still forces $p_i = |c_i|^{2\alpha}$ by compositional admissibility (B1) and unitary invariance (B2). Lemma 6 still fixes $\alpha = 1$ by the Hadamard normalisation argument, which is constructed in $d = 2$. Lemma 7 still forces iso-entropic stability. The only Lemma that requires $d \geq 3$ is Lemma 5, which fixes the trace-pairing structure for general POVMs.

In $d = 2$, the framework therefore still forces $p_i = |c_i|^2$ for projective measurements via Lemmas 4 and 6, with the trace-rule extension to general effects guaranteed by considering qubits as subsystems of larger $d \geq 3$ systems via B1. The Born rule is therefore forced in $d = 2$ as well, by a slightly different combination of Lemmas. The $d = 2$ loophole in Gleason is closed by the redundancy in the framework.

6.8 The Lüders-rule objection

The Born rule fixes outcome probabilities. What about the post-measurement state — the Lüders rule? Doesn't your framework leave that open?

Correct: the present theorem fixes the probability rule $p_i = |\langle i|\psi\rangle|^2$, not the post-measurement state $\rho_{_i} = E_i^{1/2} \rho E_i^{1/2} / \text{Tr}(E_i \rho)$. The Lüders rule is a separate question, addressed in companion work on commitment dynamics within the VERSF programme.

We note two things. First, the Lüders rule is not free given the framework: it follows from the requirement that repeated measurement of the same effect gives the same outcome with probability 1, combined with the trace-rule structure and unitarity. Second, falsification of the Lüders rule would not by itself falsify the present theorem; the two results address logically distinct features of the measurement process.

7. Relation to Prior Work

7.1 The earlier VERSF overdetermination paper

The earlier paper *Why a Fact-Producing Universe Must Satisfy the Born Rule* established the convergence of the five routes on the Born rule and analysed the residual common dependencies. It identified the kernel A0–A2 as the irreducible shared input but did not prove that the Born rule is the **unique** admissible probability rule within the framework. The present paper closes that gap.

The relationship is:

earlier paper: convergence (five routes give the same answer); present paper: inevitability (no alternative answer is admissible).

Convergence + uniqueness = inevitability. The earlier result showed that the Born rule is empirically robust across multiple derivation paths; the present result shows that the empirical robustness is the visible signature of structural necessity.

7.2 The kernel→TPB companion paper

The companion paper *Discrete Commitment and the Necessity of Tick–Bit Structure* shows that the kernel A0–A2 forces a discrete tick–bit substrate at the per-worldline level. The present paper inherits that result for Lemma 1: the U(1) phase substrate is the holonomy of the discrete tick–bit structure.

The two results are complementary:

- the kernel→TPB paper shows that the kernel forces the **temporal substrate** of physics (discrete ticks and bits along worldlines);
- the present paper shows that the kernel + derived structural conditions force both the **kinematic substrate** (complex Hilbert space, derived in Lemmas 1–3) and the **probability rule** (Born rule, derived in Lemmas 4–7) on that temporal substrate.

Together: $A0 \wedge A1 \wedge A2$ forces the temporal substrate; $A0 \wedge A1 \wedge A2 \wedge B1 \wedge B2$ forces the temporal substrate, the kinematic substrate, and the Born rule jointly. The kernel is the foundation; the derived conditions extend it to the full admissibility framework, and the present theorem now derives Hilbert kinematics rather than assuming them.

7.3 Standard reconstruction programmes

The present result differs from existing reconstruction programmes (Gleason 1957, Deutsch–Wallace 1999/2007, Zurek 2003/2005, Hardy 2001, Masanes–Müller 2011) in two structural ways.

First, modal status. Existing reconstructions establish that *some* sufficient axiom set forces the Born rule: each derivation succeeds, but each rests on an axiom whose status is no less contested than the rule it derives. Deutsch–Wallace requires a decision-theoretic rationality postulate; Zurek requires envariance; Hardy requires a continuity-of-states axiom. The contested axiom does the work, and the philosophical status of the contested axiom inherits the strangeness that motivated the reconstruction in the first place.

The present result establishes that the Born rule is forced by a structural framework whose conditions are *individually* non-contentious — A0, A1, A2 are kernel principles of any observer-invariant, finite-resource, fact-producing theory; B1, B2 are derived structural requirements of having compositional dynamics at all. The contestation moves from the rule-deriving axiom to the entire framework, which is harder to attack because each component is more deeply motivated.

Second, unification level. Existing reconstructions are single-route: one set of axioms, one derivation, one result. The present result is multi-route unified: five sets of axioms, five derivations, one result, plus a uniqueness theorem showing that the result is the unique fixed point of the joint constraints. This is structurally richer than any single-route reconstruction and provides a different kind of foundational support: not just that the rule *can* be derived, but that the rule is the unique admissible measure within a multiply-constrained framework.

8. Falsification Conditions

A foundational result of this kind must commit to empirical conditions for its defeat. We list four.

(F1) Demonstration of admissible non-quadratic probability. If an experiment showed that probability assignments in a fact-producing context took the form $p_i = |c_i|^p$ for some $p \neq 2$, while satisfying all of A0, A1, A2, B1, and B2, the Admissible Probability Uniqueness Theorem would be falsified. This would require a measurement context in which compositional admissibility (B1) and unitary normalisation (B2) hold but Hadamard-type counterexamples to $p \neq 2$ fail to manifest. Such an experiment would also falsify the existing single-route Born-rule reconstructions; the present result would simply be one more casualty.

(F2) Demonstration of admissible contextual probability. If an experiment showed that probability assignments depend on the full POVM context in a way that survives all of the admissibility constraints — i.e. context-dependence that is observer-invariant under A0 and compositionally consistent under B1 — the contextuality-loophole closure of §6.5 would be falsified.

(F3) Detection of non-iso-entropic deviations. Route V predicts deviations from the Born rule of the form

$$p_i \propto |c_i|^2 \cdot \exp(-\lambda \Delta S_i)$$

in the non-iso-entropic regime. Detection of such deviations would *not* falsify the present result — the Born rule is the iso-entropic limit, and the broader law contains it as a special case — but would *confirm* the structural framework while extending it. Failure to detect such deviations in regimes where ΔS_i is engineered to vary across outcomes would indicate that Route V's stability claim holds more strongly than expected, which would be a positive result for the framework.

(F4) Inconsistency in the kernel→TPB chain. Lemma 1 relies on the kernel→TPB result for the U(1) substrate. If the kernel→TPB derivation were shown to fail — e.g., if $A0 \wedge A1 \wedge A2$ admitted models without the discrete tick-bit substrate — then Lemma 1 would lose its foundation, and the present theorem would require a different route to phase structure. This is a falsification of a presupposition rather than of the theorem itself, but is listed for completeness.

(F1) and (F2) are direct falsification conditions. (F3) is a discrimination condition between the strict Born rule and the broader entropic-unfolding law that contains it. (F4) is a foundational dependency.

9. Limitations and Scope

We are explicit about what this paper does and does not establish.

(L1) Inheritance of the kernel. The result rests on the kernel A0–A2 and on the derived conditions B1–B2. The kernel is foundational to the broader VERSF programme; the derived conditions are themselves justified within the programme. The result is therefore conditional on the broader VERSF foundational commitments, not free-standing. A reader who rejects the kernel rejects the framework, but the rejection must engage with the kernel principles individually rather than with the rule they jointly force.

(L2) The five routes are themselves results of the broader programme. Each of Routes I–V is established in dedicated companion papers. The unification depends on the prior establishment of the routes; the present paper does not re-derive them. The strengthening comes from showing that they are constraints on a single object, not from new derivations.

(L3) Step 5 and the entropic regime. The Born rule emerges from Route V as the iso-entropic limit. The non-iso-entropic regime predicts departures from the Born rule that are tested by experiments specifically designed for non-iso-entropic conditions; the present result establishes only that the Born rule is the unique admissible measure in the iso-entropic regime that characterises standard quantum experiments. The deeper law $p_i \propto |c_i|^2 \cdot \exp(-\lambda \Delta S_i)$ is admissible in the broader sense; the strict Born rule is admissible in the iso-entropic limit.

(L4) Operational vs model-theoretic independence. As in the earlier overdetermination paper, the independence of the five routes is established at the operational level (distinct primitives, distinct proofs) rather than at the strict model-theoretic level (no shared models). Strict model-theoretic independence remains an open technical question, though the operational independence is sufficient to defeat shared-error explanations of the convergence.

(L5) Hilbert kinematics now derived from \mathfrak{A} . Earlier drafts of this work invoked Hilbert structure as an external presupposition at Lemma 5 (the Gleason–Busch step). The substrate lemmas (Lemmas 1–3) close that gap: complex Hilbert structure is now derived from \mathfrak{A} itself, not inherited. Two presuppositions remain: dimension $d \geq 3$ for Gleason–Busch (the $d = 2$ case is handled redundantly by Lemmas 4 and 6, see §6.7), and the existence of a continuous family of admissible reversible dynamics (used in Lemma 3 to invoke Haar averaging). Both are operationally well-motivated within the kernel and do not constitute external Hilbert assumptions.

(L6) The Lüders rule. The theorem fixes the probability rule, not the post-measurement state. The Lüders rule is a separate question (§6.8).

(L7) "Inevitability" is framework-relative. The precise claim of the theorem is: **the Born rule is the unique admissible probability measure within any fact-producing framework satisfying \mathfrak{A} .** This is a framework-relative uniqueness result, not a claim that the Born rule holds in all logically possible universes. A reader who rejects the kernel A0–A2 or the derived structural conditions B1–B2 is not committed to the Born rule by the present argument. The strength of the argument lies elsewhere: in the difficulty of rejecting any individual component of \mathfrak{A} while retaining a coherent fact-producing physical theory. We do not claim that \mathfrak{A} is the only possible framework — we claim that within it, the Born rule is uniquely forced, and that each component of \mathfrak{A} is independently motivated to a degree that makes its rejection costly.

10. Conclusion

This paper began as a unification of five independent VERSF derivations of the Born rule. In closing the gaps that earlier drafts left open — the kernel-circularity of B1 and B2, the inheritance of Hilbert structure, the field-selection question, the compactness of the symmetry group — it has become something larger. The substrate lemmas (1–3) derive complex Hilbert kinematics from \mathfrak{A} itself; the rule lemmas (4–7) derive the Born rule on that derived substrate. Together they exhibit a chain that runs from the conditions for fact-production to the full kinematic and probabilistic structure of quantum mechanics:

Facts ($A0 \wedge A1 \wedge A2$) \downarrow Admissibility framework ($\mathfrak{A} = A0 \wedge A1 \wedge A2 \wedge B1 \wedge B2$) \downarrow Complex Hilbert state space (Lemmas 1–3: $U(1)$ phase, local tomography, compact symmetry \rightarrow invariant inner product) \downarrow Unitary dynamics ($B2 + \text{Lemma 3}$) \downarrow Born rule $p_i = |\langle i|\psi\rangle|^2$ (Lemmas 4–7: bilinear core, trace structure, $p = 2$, iso-entropic stability).

Each layer is forced by the conjunction of layers above. None is a free postulate. The strongest defensible claim of this work is therefore not just about probability, but about the architecture itself:

Quantum kinematics (complex Hilbert space, unitary dynamics) and the Born rule jointly emerge as the unique admissible structure consistent with $A0 \wedge A1 \wedge A2 \wedge B1 \wedge B2$.

The seven-lemma proof is the technical content. Lemma 2 — the No-Alternative Theorem for Composition — is the linchpin: it forces \mathbb{C} as the amplitude field by exhausting the three failure modes of non-tomographic composition (hidden global parameters, non-reconstructible states, non-factorising marginals), each of which contradicts at least one component of \mathfrak{A} . The five existing Routes converge on the Born rule because they are five projections of this single architectural structure; the architecture itself is forced by the kernel together with two derived structural conditions whose individual rejection produces operationally absurd theories (observer-dependent outcome frequencies, premature commitment, spontaneous fact generation).

The killer line:

Any deviation from quadratic probability either fails to be a fact-producing, composable, observer-invariant physical theory at all, or violates at least one admissibility constraint in \mathfrak{A} .

The result reframes the status of the Born rule. It is no longer a postulate, no longer one of several possible rules, no longer merely the answer that five independent routes happen to converge on. It is **the unique admissible probability measure within any fact-producing framework satisfying \mathfrak{A}** . Hilbert kinematics, local tomography, \mathbb{C} as the amplitude field, compactness of the symmetry group, the trace pairing, the exponent $p = 2$, and iso-entropic stability are all derived consequences of \mathfrak{A} , not external presuppositions. The contestation, if any, must occur at the layer of the kernel and the derived structural conditions — at the level of *what counts as a fact-producing physics* — rather than at the kinematic or probabilistic layers.

Final statement

Within VERSF, the Born rule is not merely overdetermined. It is the unique admissible probability measure compatible with finite distinguishability, irreversible fact formation, compositional structure, and observer-invariant physics — and the complex Hilbert kinematics on which it acts are themselves derived from these conditions, not assumed. Any alternative probability rule would require a framework that cannot preserve facts, cannot compose systems, or cannot maintain observer-invariant consistency. Quantum mechanics, in both its kinematic and probabilistic structure, is therefore the unique admissible architecture for a fact-producing, composable, observer-invariant physical theory.

Notation Summary

Symbol	Meaning
\mathfrak{A}	the admissibility framework: $A0 \wedge A1 \wedge A2 \wedge B1 \wedge B2$
A0	observer-invariant distinguishability (kernel)
A1	finite admissibility (kernel)
A2	irreversible commitment (kernel)
B1	compositional admissibility (derived structural condition)
B2	reversible pre-commitment dynamics (derived structural condition)
\mathcal{C}	commitment map $\mathcal{A} \rightarrow \mathcal{O}$
\mathcal{G}	group of admissible observer transformations
\mathcal{S}	a bounded subsystem
\mathcal{A}	set of admissible alternatives
\mathcal{O}	set of committed outcomes
$c_i = \langle i$	$\psi \rangle$
p_i	probability assigned to outcome i
ρ	quantum state (density operator)
E_i	effect (POVM element) for outcome i
ΔS_i	entropic action associated with outcome i
λ	inverse-temperature-like coefficient in entropic action
α	exponent in $p_i =$
Routes I–V	the five existing VERSF derivations of the Born rule
Lemmas 1–7	the seven constraints in the uniqueness theorem proof: 1–3 substrate (Hilbert structure), 4–7 probability rule
\mathcal{K}	the group of admissible reversible pre-commitment transformations on a bounded subsystem
$\mathcal{H}(\mathcal{S})$	the admissible state space (a finite-dimensional complex Hilbert space) on subsystem \mathcal{S} ; derived in Lemmas 1–3

Selected References

- Born, M. (1926). Zur Quantenmechanik der Stoßvorgänge. *Z. Phys.* 37, 863–867.
- Gleason, A. M. (1957). Measures on the closed subspaces of a Hilbert space. *J. Math. Mech.* 6, 885–893.
- Busch, P. (2003). Quantum states and generalized observables: a simple proof of Gleason's theorem. *Phys. Rev. Lett.* 91, 120403.
- Caves, C. M., Fuchs, C. A., Manne, K. K., Renes, J. M. (2004). Gleason-type derivations of the quantum probability rule for generalized measurements. *Found. Phys.* 34, 193–209.
- Deutsch, D. (1999). Quantum theory of probability and decisions. *Proc. R. Soc. A* 455, 3129–3137.
- Wallace, D. (2007). Quantum probability from subjective likelihood: improving on Deutsch's proof of the probability rule. *Stud. Hist. Philos. Mod. Phys.* 38, 311–332.
- Zurek, W. H. (2005). Probabilities from entanglement, Born's rule from envariance. *Phys. Rev. A* 71, 052105.
- Hardy, L. (2001). Quantum theory from five reasonable axioms. arXiv:quant-ph/0101012.
- Masanes, L., Müller, M. P. (2011). A derivation of quantum theory from physical requirements. *New J. Phys.* 13, 063001.
- Jaynes, E. T. (1957). Information theory and statistical mechanics. *Phys. Rev.* 106, 620–630.
- Pressé, S., Ghosh, K., Lee, J., Dill, K. A. (2013). Principles of maximum entropy and maximum caliber in statistical physics. *Rev. Mod. Phys.* 85, 1115–1141.
- Lüders, G. (1951). Über die Zustandsänderung durch den Messprozeß. *Annalen der Physik* 8, 322–328.
- Landauer, R. (1961). Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development* 5, 183–191.
- Taylor, K. (VERSF programme): *Why a Fact-Producing Universe Must Satisfy the Born Rule* (overdetermination paper); *Discrete Commitment and the Necessity of Tick–Bit Structure* (kernel→TPB paper); *Foundations of Tick–Bit Dynamics*; *The Double Square Rule and Bilinear Probability*; *Physical Admissibility Framework: Eliminating $p \neq 2$* ; *Maximum-Caliber Entropic Unfolding and the Born Rule*. versf-eos.com.