Probability as Consistency: Eliminating the Born Rule as an Independent Postulate

A Single-Source Derivation in the VERSF Framework

Keith Taylor *VERSF Theoretical Physics Programme*

General-Reader Abstract

What *is* a probability? Outside physics, there are roughly three answers in circulation. The **frequentist** view says probability is the long-run frequency of an outcome over many repetitions of an experiment — flipping a coin a thousand times and counting heads. The **Bayesian** view says probability is a degree of belief held by a rational agent, updated as new evidence arrives. The **propensity** view says probability is an objective tendency built into a physical situation, prior to any agent or any sequence of trials. Each captures something genuine; none is universally accepted. In ordinary life and in classical statistical mechanics, these readings can usually be reconciled — probabilities track our ignorance about underlying configurations we could in principle uncover.

Quantum mechanics breaks this comfortable picture. When a quantum measurement is performed, the outcome is unpredictable in a way that does not appear to reduce to ignorance: the universe itself seems not to have decided which outcome will occur until the measurement happens. The *probabilities* of different outcomes nevertheless follow a precise mathematical rule — the squared modulus of a complex-valued amplitude, known as the **Born rule**. In the textbook formulation this rule is simply assumed. It is one of the foundational axioms of the theory, sitting alongside the rules for how quantum states evolve, and it is logically separate from them: nothing in the rest of the formalism *makes* the rule come out this way. For nearly a century, physicists have asked whether this is a brute fact about nature, or whether the Born rule can be derived from something more fundamental. Several attempts exist — Gleason's theorem, decision-theoretic arguments, environment-induced symmetry arguments — but each replaces the Born postulate with a different one whose physical status is itself contested.

This paper argues that within VERSF — a theoretical framework in which all of physics is built from a single quantity, the density of irreversible facts produced by the universe — the Born rule is not an extra ingredient at all. It is forced by the structure of the framework itself.

The argument is straightforward in outline. Probability is reinterpreted as a measure of *consistency*: how compatible a particular candidate outcome is with the existing record of facts the universe has already produced. We show that four natural requirements on this consistency measure — that it be non-negative, that it not depend on counterfactual outcomes that did not occur, that it be invariant under the reversible dynamics that precede a measurement, and that it

factorise across causally disjoint parts of the universe — uniquely determine it to be the Born rule. The same structure also forces an information-theoretic reading in which probability and entropy turn out to be two faces of one quantity. We further show that *no alternative probability rule is possible* within the framework: any candidate either depends on hidden data the framework forbids, or violates one of the four requirements.

The result removes one of the standing redundancies in the foundations of quantum mechanics. Probability is neither a frequency over imagined repetitions, nor a degree of belief, nor an unexplained propensity. It is the geometry of how facts compose. The Born rule is not a separate postulate of the universe — it is what that geometry permits, and nothing else is permitted at all.

Table of Contents

Abstract

In the Void Energy–Regulated Space Framework (VERSF), every observable is a functional of one scalar field — the committed record density **ρ(x,t)** — encoding the cumulative distribution of irreversible commitment events. Prior work in the programme establishes the structural origin of the pre-commitment state space, unitary evolution, measurement, geometry, and the fundamental constants. The status of *probability*, however, has so far been left implicit.

This paper closes that gap. We define a candidate commitment C_i as **consistent** with ρ to the degree that admissible observables are invariant under admissible extensions $\rho \to \rho + \delta\rho\_\{C_i\}$

preserving reversible (unitary) pre-commitment structure, and we show that probability is the emergent measure quantifying this invariance over the admissible extension set $\mathcal{C}$. The argument rests on four constraints — non-negativity, non-contextuality (forced by the single-source theorem), unitary invariance (forced by the derived structure of pre-commitment dynamics), and compositional multiplicativity (forced, via Proposition 1, by disjoint-sector decomposition of ρ). Theorem 1 proves that

$$\mathbf{p_i = |\langle i|\psi\rangle|^2 \equiv |c_i|^2}$$

is the unique consistency measure compatible with these constraints, in every Hilbert dimension ≥ 2. Theorem 2 strengthens this to a no-go statement: no admissible probability structure outside the Born rule can exist within VERSF. Lemma 1 then derives the extension-information functional $\Delta I_i = -\ln |c_i|^2$ as the unique extensive measure of extension cost, so that Corollary 1 obtains the Gibbs form

$$\mathbf{p_i = \exp(-\Delta I_i)}$$

without independent input. The Born rule is therefore not an independent postulate of VERSF but the only consistency measure compatible with the structure of the record. In a single identity it unifies probability, entropy, information geometry, and κ-field coupling, collapsing one of the standing axiomatic redundancies of quantum mechanics and completing the single-source identification

$$\boldsymbol{\mathcal{O} = \mathcal{F}[\rho]}$$

for probabilistic observables.

---

1. Introduction

In the standard formulation of quantum mechanics, the Born rule

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

is an independent postulate, logically separate from the kinematical postulate (states as rays in a complex Hilbert space) and from the dynamical postulate (unitary evolution). This separation is the source of considerable foundational discomfort: probability appears as an additional ingredient, neither implied by the rest of the formalism nor visibly compatible with it.

Several derivations have attempted to close the gap — Gleason's theorem, the Deutsch–Wallace decision-theoretic programme, Zurek's envariance argument, and frequentist proposals of various kinds. Each succeeds, but only under additional assumptions whose physical status is itself contested.

VERSF takes a different starting point. Its core claim is the **single-source theorem**: every physical observable is a functional of one scalar field, the committed record density,

$$\mathcal{O} = \boldsymbol{\mathcal{F}}[\rho(x,t)],$$

with no independent fact-producing structure. In this architecture, an *independent* probability postulate would be a conspicuous anomaly — a degree of freedom not carried by ρ. The present paper shows that no such anomaly is needed. Probability, like geometry, dynamics, and the fundamental constants, is fixed by the admissibility structure of the record.

The argument proceeds in three steps. Section 2 sets out the candidate-extension structure of ρ and the associated weighting functional W. Section 3 imposes four admissibility constraints, each grounded in prior VERSF results rather than introduced for the present argument; the most contentious of these — compositional multiplicativity — is upgraded to a theorem (Proposition 1) about ρ-decomposition. Section 4 proves that these constraints uniquely determine W to be the squared modulus of the pre-commitment amplitude (Theorem 1), strengthens this to a no-go statement (Theorem 2), and identifies the resulting probability measure with a Gibbs distribution over a uniquely-determined extension-information functional (Lemma 1, Corollary 1). Sections 5–7 interpret the result, situate it relative to the existing literature, and set out scope.

This paper is **complementary** to the earlier *Born Rule as Entropic Unfolding* result of the programme. That work derives $p_i = |c_i|^2$ from informational unfolding of pre-commitment structure; the present work derives the same rule from the consistency structure of the record itself. Corollary 1 shows that the two routes converge on the same expression because they are two faces of the same single-source architecture.

---

2. Framework

2.1 The committed record

A **commitment event** is an irreversible transition from a structured set of distinguishable possibilities to a single realised outcome. The committed record density

$$\rho(x,t)$$

encodes the spatiotemporal distribution of such events. ρ is the carrier of all distinguishability in the theory: a state of affairs *is* a configuration of ρ. There are no fact-producing degrees of freedom orthogonal to ρ. This is the content of the single-source theorem.

2.2 Pre-commitment algebra

Reversible pre-commitment structure has been derived elsewhere in the programme, from the admissibility axioms (finite distinguishability, irreversible commitment, compositional closure) together with the requirement that the universe support interference without producing facts. The derivation forces a complex Hilbert space $\mathscr{H}$ of finite local dimension, with the minimal local algebra being $\mathbb{C}^4$. For the present argument we need only the following:

- **(P1)** Pre-commitment states are rays in a complex Hilbert space $\mathscr{H}$, dim $\mathscr{H} \geq 2$.
- **(P2)** Reversible pre-commitment dynamics act by unitary operators $U \in \mathcal{U}(\mathscr{H})$.
- **(P3)** Measurement contexts correspond to orthonormal resolutions of the identity $\{|i\rangle\}$, with $\sum_i |i\rangle\langle i| = \mathbb{1}$.

A pre-commitment state expanded in a measurement basis takes the form

$|\psi\rangle = \sum_i c_i |i\rangle$, $c_i \in \mathbb{C}$, $\sum_i |c_i|^2 = 1$.

(P1)–(P3) are inputs to this paper but outputs of prior work; they are not assumed here on their own merits. The argument that follows derives one further structural result — the unique probability measure on (P3) — from the same admissibility skeleton that produced (P1)–(P3).

2.3 Candidate extensions

Given a state ρ and a measurement context $\{|i\rangle\}_{i=1..d}$, the **candidate-extension set** is

$\mathcal{C}(\rho, \{|i\rangle\}) = \{ C_1, C_2, \ldots , C_d \}$,

where each C_i corresponds to the outcome $|i\rangle$ being committed and the resulting record being $\rho + \delta\rho_{C_i}$. Admissibility — finite distinguishability, irreversibility, and compositional closure — restricts which C_i remain in $\mathcal{C}$.

2.4 Consistency and the weighting functional

We make the central concept precise.

Definition (Consistency). A candidate extension $C_i \in \mathcal{C}$ is **consistent** with ρ to the degree that *admissible observables* are invariant under *admissible extensions* $\rho \to \rho + \delta\rho_{C_i}$ *preserving reversible (unitary) pre-commitment structure*. Three qualifications are intended explicitly by this phrasing:

- **Admissible observables.** The invariance condition applies only to observable functionals $\mathcal{F}[\rho]$ permitted by the single-source theorem; non-admissible functionals are not in scope.
- **Admissible extensions.** Only candidate extensions $C_i \in \mathcal{C}$ — those satisfying finite distinguishability, irreversibility, and compositional closure — are evaluated; inadmissible extensions carry no consistency weight by definition.
- **Reversible pre-commitment structure.** The invariance is evaluated within the unitary pre-commitment regime (P2); regimes that destroy reversibility (e.g. via non-unitary fact-production processes outside C_i itself) are excluded from the comparison.

Consistency is therefore a structural invariance condition on a closed sub-class of observables, against a closed sub-class of extensions, within a closed sub-class of dynamics. There is no remaining ambiguity in what "invariance under extension" means.

The **consistency weighting** is the functional

$W : \mathcal{H} \times \{\text{basis ray}\} \to \mathbb{R}_{\geq 0}, (|\psi\rangle, |i\rangle) \mapsto W_i(|\psi\rangle),$

quantifying this degree of invariance for each candidate extension. Probabilities are obtained by normalisation,

$p_i = W_i / \sum_j W_j.$

The question this paper answers is sharp: *what does admissibility force W to be?*

---

3. Admissibility constraints on W

We impose four constraints. Each is grounded in independent VERSF structure rather than postulated for the present argument.

(A1) Non-negativity

$W_i(|\psi\rangle) \geq 0$ for all admissible C_i.

Justification. W measures structural compatibility between an admissible candidate extension and the existing record. Negative compatibility has no physical reading in a framework where only realised commitments count as facts.

(A2) Non-contextuality

$W_i(|\psi\rangle)$ depends only on $|i\rangle$ and $|\psi\rangle$, not on the other elements of the orthonormal resolution containing $|i\rangle$.

Justification. The single-source theorem forbids hidden contextual dependence. The consistency of C_i with ρ cannot depend on counterfactual basis elements that are not committed; if it did, the probabilistic observable p_i would not be a functional of ρ alone, in direct contradiction of $\mathcal{O} = \mathcal{F}[\rho]$. Equivalently: if two measurement contexts share the outcome ray $|i\rangle$ but differ on its orthogonal complement, the consistency weighting for C_i must agree across them.

This is the VERSF analogue of the non-contextuality assumption that Gleason's theorem makes on frame functions, but here it is **derived** from the single-source theorem rather than postulated.

(A3) Unitary invariance

$W_i(U|\psi\rangle; U|i\rangle) = W_i(|\psi\rangle; |i\rangle)$ for all $U \in \mathcal{U}(\mathcal{H})$.

Justification. Reversible pre-commitment dynamics produce no facts — this is the defining property separating pre-commitment evolution from commitment, and the same property invoked in the consistency definition's "preserving reversible structure" clause. They cannot, therefore, alter the consistency weighting attached to any candidate extension. The unitary structure of these dynamics is itself a derived feature of admissibility (P2), not assumed here on independent grounds.

(A4) Compositional multiplicativity

For independent subsystems A and B in a joint product state $|\psi_A\rangle \otimes |\psi_B\rangle$, with joint outcome (i, j) corresponding to $|i\rangle_A \otimes |j\rangle_B$,

$$W_{(i,j)}(|\psi_A\rangle \otimes |\psi_B\rangle) = W_i(|\psi_A\rangle) \cdot W_j(|\psi_B\rangle).$$

Justification. (A4) is not an independent axiom but a corollary of single-source structure applied to disjoint sectors of ρ. We make this precise.

Proposition 1 (Disjoint-Sector Factorisation)

Let the committed record decompose as

$$\boldsymbol{\rho = \rho_A \oplus \rho_B}$$

into causally disjoint sectors — that is, sectors such that no admissible observable functional $\mathcal{F}[\rho]$ contains cross-terms between ρ_A and ρ_B. Then:

(i) The admissible extension set factorises set-theoretically:

$$\boldsymbol{\mathcal{C}(\rho) = \mathcal{C}(\rho_A) \times \mathcal{C}(\rho_B).}$$

(ii) Every admissible consistency weighting W is multiplicative across sectors:

$$\boldsymbol{W(C_{(i,j)}) = W^A(C_i) \cdot W^B(C_j),}$$

up to overall normalisation.

Proof.

Part (i). A joint candidate extension $\delta\rho_{C_{(i,j)}} = \delta\rho^A_{C_i} \oplus \delta\rho^B_{C_j}$ is admissible iff its sectoral components are each admissible: by disjointness, no cross-sector consistency condition exists, and admissibility evaluates independently in each sector. Conversely, every pair $(C_i, C_j) \in \mathcal{C}_A \times \mathcal{C}_B$ yields an admissible joint extension. Hence $\mathcal{C}(\rho) = \mathcal{C}_A \times \mathcal{C}_B$.

Part (ii). By the single-source theorem $\mathcal{O} = \boldsymbol{\mathcal{F}}[\rho]$, the weighting W is itself a functional of ρ and the candidate extension. Disjointness forbids W from depending on cross-terms between sectors; $W(C_{(i,j)})$ must therefore be a function of the marginal weights alone:

$$W(C_{(i,j)}) = h(W^A(C_i), W^B(C_j))$$

for some $h : \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$.

Marginal compatibility now fixes h. The marginal probability assigned to $C_i \in \mathcal{C}\_A$ by the joint weighting,

$$p^A(C_i) = \sum_j W(C_{(i,j)}) / \sum_{\{k,\ell\}} W(C_{(k,\ell)}),$$

must coincide with the marginal probability computed within sector A in isolation,

$$p^A(C_i) = W^A(C_i) / \sum_k W^A(C_k);$$

otherwise the consistency structure of A would depend on the existence or content of B, contradicting the single-source theorem applied to the disjoint decomposition.

Imposing this for all admissible distributions on B forces

$$\sum_j h(W^A(C_i), W^B(C_j)) = W^A(C_i) \cdot \kappa_B,$$

where κ_B is some constant depending only on the B-data. Since this must hold for every i and every admissible W^A, h must be linear in its first argument; by symmetry of A and B, linear in its second. Combined with $h \geq 0$ and continuity (inherited from (A3)), this forces

$$h(x, y) = \kappa \cdot x \cdot y, \; \kappa > 0.$$

Any non-multiplicative term in h would couple the marginal of A to the partition function of B, which is forbidden by the single-source theorem applied to disjoint sectors. Absorbing κ into normalisation gives $W(C_{(i,j)}) = W^A(C_i) \cdot W^B(C_j)$. ■

Status. Proposition 1 supplies (A4) as a theorem about ρ rather than an assumption about state vectors. The objection that (A4) imports tensor-product intuition from quantum mechanics is therefore void: multiplicativity is forced by the single-source theorem applied to causally disjoint sectors of the record, with no reference to the Hilbert-space tensor product.

These constraints do not float free: (A1) is the meaning of "consistency"; (A2) is forced by single-source structure; (A3) is inherited from the derived unitary character of pre-commitment dynamics; (A4) is forced by Proposition 1 from disjoint-sector decomposition of ρ.

---

4. The uniqueness theorems

4.1 The consistency-weighting theorem

We now show that (A1)–(A4), together with the additivity of probability over disjoint outcomes, fix W up to an overall normalisation.

Theorem 1 (Uniqueness of the consistency weighting)

Let dim $\mathscr{H} \geq 2$. Any functional $W : \mathscr{H} \times \{\text{basis ray}\} \to \mathbb{R}_{\geq 0}$ satisfying (A1)–(A4) and assigning, after normalisation, a probability measure on the outcome space takes the form

$$W_i(|\psi\rangle) = \kappa\, |\langle i|\psi\rangle|^2,\ \kappa > 0,$$

and after normalisation

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

Proof

Step 1 — Reduction to a function of $|c_i|^2$.

By (A2), $W_i(|\psi\rangle)$ depends only on $|i\rangle$ and $|\psi\rangle$. By (A3), it is invariant under simultaneous unitary action on the pair $(|\psi\rangle, |i\rangle)$. The only continuous invariant of a pair of rays in $\mathscr{H}$ under the diagonal unitary action $\mathcal{U}(\mathscr{H})$ is the squared overlap $|\langle i|\psi\rangle|^2 \in [0,1]$. Hence there exists a function

$$f : [0,1] \to \mathbb{R}_{\geq 0},\ W_i(|\psi\rangle) = f(|\langle i|\psi\rangle|^2) = f(|c_i|^2).$$

The continuity of f (used in the next step) is inherited: $\mathcal{U}(\mathscr{H})$ acts continuously on $\mathscr{H}$, and any W invariant under (A3) varies continuously with the amplitudes.

Step 2 — Multiplicativity fixes f up to an exponent.

Apply (A4) (now grounded in Proposition 1) to product states of the form

$$|\psi_A\rangle = \sqrt{a}\, |i_A\rangle + \sqrt{(1-a)}\, |i_A^\perp\rangle,\ |\psi_B\rangle = \sqrt{b}\, |i_B\rangle + \sqrt{(1-b)}\, |i_B^\perp\rangle,$$

with $a, b \in [0,1]$. The joint amplitude on (i_A, i_B) is $\sqrt{(ab)}$, so its squared modulus is ab. (A4) then requires

$$f(ab) = f(a) \cdot f(b) \text{ for all } a, b \in [0,1].$$

This is the multiplicative Cauchy equation on [0,1]. Combined with non-negativity (A1) and continuity, the standard classification of its solutions gives either $f \equiv 0$, $f \equiv 1$ (a constant solution), or

$f(x) = x^\alpha$, $\alpha > 0$.

The trivial solution $f \equiv 0$ is excluded by non-degenerate normalisation. The constant solution $f \equiv 1$ is excluded by the same: in any dim ≥ 2, $\sum_i f(|c_i|^2) \neq 1$ in general for $f \equiv 1$. Hence $f(x) = x^\alpha$ for some $\alpha > 0$, and in particular $f(0) = 0$.

Step 3 — Coarse-graining additivity fixes α = 1.

Fix a state $|\psi\rangle$ and an orthonormal resolution $\{|i\rangle\}_{i=1..d}$. Consider replacing the resolution by another that agrees on $|3\rangle, |4\rangle, \ldots, |d\rangle$ but uses an alternative orthonormal pair $(|+\rangle, |-\rangle)$ within span$\{|1\rangle, |2\rangle\}$. By (A2), the weighting on $|3\rangle, \ldots, |d\rangle$ is unchanged. By normalisation,

$f(|c_1|^2) + f(|c_2|^2) = f(|c_+|^2) + f(|c_-|^2)$, $(*)$

where (c_+, c_-) are the components of $|\psi\rangle$ in the alternative pair. Equation $(*)$ shows that the sum on the left depends only on the projection of $|\psi\rangle$ onto span$\{|1\rangle, |2\rangle\}$, not on the choice of orthonormal pair within that subspace.

By (A3), any unitary fixing span$\{|1\rangle, |2\rangle\}$ pointwise leaves c_1, c_2 unchanged; any unitary acting within that subspace only relabels the orthonormal pair, which by the previous sentence leaves the sum invariant. Hence $f(|c_1|^2) + f(|c_2|^2)$ depends only on the squared norm of the projection,

$\|P_{12} |\psi\rangle\|^2 = |c_1|^2 + |c_2|^2$.

Define $g(s) := f(|c_1|^2) + f(|c_2|^2)$ where $|c_1|^2 + |c_2|^2 = s$. By varying $(|c_1|^2, |c_2|^2)$ over all decompositions of $s \in [0,1]$ this is well-defined.

Now set $|c_2|^2 = 0$. Then $s = |c_1|^2$, and $g(s) = f(s) + f(0) = f(s)$ (using $f(0) = 0$ from Step 2). Combining this with the general expression,

$f(x) + f(y) = f(x + y)$ for all $x, y \geq 0$ with $x + y \leq 1$.

This is the additive Cauchy equation on $[0,1]$, and combined with $f(x) = x^\alpha$ from Step 2 it requires

$(x + y)^\alpha = x^\alpha + y^\alpha$.

Setting $x = y = t$ with $2t \leq 1$ gives $(2t)^\alpha = 2t^\alpha$, hence $2^\alpha = 2$, so **$\alpha = 1$**.

Step 4 — Normalisation.

With $f(x) = x$, the weighting is $W_i(|\psi\rangle) = \kappa\, |c_i|^2$ for some $\kappa > 0$. Normalisation gives

$p_i = |c_i|^2 / \sum_j |c_j|^2 = |c_i|^2$,

since $|\psi\rangle$ is normalised. ■

Remarks on the proof

- **Validity in dimension 2.** The argument applies in dim $\mathscr{H} = 2$, where Gleason's theorem fails. The crucial replacement is (A4): compositional multiplicativity does the work in dim 2 that frame-function additivity does at dim ≥ 3 in Gleason's argument. Via Proposition 1, (A4) has independent physical content (single-source structure on disjoint sectors of ρ) and is not an ad hoc strengthening introduced to cover the qubit case.
- **No probabilistic input beyond (A1)–(A4).** Step 3 invokes only the additivity of probability over disjoint outcomes — that is, the *definition* of a probability measure — together with (A2) and (A3). No frequency, decision-theoretic, or self-locating axiom is required.
- **Multiplicativity and additivity together are necessary.** Multiplicativity alone gives $f(x) = x^\alpha$ for any $\alpha > 0$. Additivity alone gives $f(x) = \alpha x$. Their conjunction forces $f(x) = x$ exactly.
- **f(0) = 0 was needed.** Step 2 establishes this directly from the multiplicative form (and the rejection of the constant solution by normalisation), so the use of $f(0) = 0$ in Step 3 introduces no new assumption.

4.2 The no-go theorem

Theorem 1 shows that the Born rule is the unique consistency weighting compatible with admissibility. We now strengthen this to the assertion that *no other admissible probability structure can exist at all* within VERSF.

Theorem 2 (No Independent Probability Postulate)

Within the VERSF admissibility class, no probability structure on commitment outcomes can be defined independently of, or as an alternative to, the Born weighting $p_i = |c_i|^2$.

Proof

Suppose an alternative admissible probability assignment p'_i exists, distinct from the Born weighting. Two cases exhaust the possibilities.

Case 1 — p' depends on data outside ρ. Suppose

$p'_i = \mathcal{F}$_p[ρ, X],

with X non-trivial and not reducible to a functional of ρ. Then X encodes distinguishability that is not carried by ρ. But the single-source theorem $\mathcal{O} = \mathcal{F}[\rho]$ asserts that all admissible observables are functionals of ρ alone, so X can carry no admissible distinguishability. Either X is a trivial constant — in which case it does not enter — or $\mathcal{F}$_p is not an admissible observable. The supposition fails.

Case 2 — p' is a functional of ρ alone but differs from the Born rule. Suppose

$p'_i = f(|c_i|^2)$, $f \neq$ identity.

By admissibility, p' must satisfy:

- non-negativity (else p' is not a probability),
- non-contextuality (else p' is not a functional of ρ alone, contradicting Case 1's exclusion),
- unitary invariance (else it depends on pre-commitment dynamics that produce no facts, in violation of (P2)),
- compositional multiplicativity (else it violates Proposition 1 on disjoint sectors).

These are precisely (A1)–(A4). Theorem 1 then forces f = identity, contradicting the supposition.

Both cases are excluded. Therefore no admissible probability structure exists outside

$p_i = |c_i|^2$. ■

The result is sharp: the Born rule is not merely *consistent* with VERSF admissibility — it is the *only* probability structure compatible with it. Any candidate alternative either violates the single-source theorem (Case 1) or one of the four admissibility constraints (Case 2). The space of admissible probability structures within VERSF contains exactly one element.

4.3 The information-theoretic form

The Born weighting admits an equivalent expression as a Gibbs measure. To make this precise we first determine the unique extensive measure of extension cost.

Lemma 1 (Uniqueness of the extension-information functional)

Let $\Delta I : (0, 1] \to \mathbb{R}_{\geq 0}$ be a measure of the information cost of extending ρ by a candidate C_i with squared amplitude $|c_i|^2 = x$, satisfying:

(L1) Additivity under composition: $\Delta I(ab) = \Delta I(a) + \Delta I(b)$ for a, b $\in (0,1]$.

(L2) Continuity on (0, 1].

(L3) Monotonicity: ΔI is non-increasing in x, with $\Delta I(1) = 0$.

Then

$$\mathbf{\Delta I(x) = -\lambda \ln x}$$

for some constant λ > 0.

Proof

Let $g(t) := \Delta I(e^{-t})$ for $t \in [0, \infty)$. Then by (L1),

$g(s + t) = \Delta I(e^{-s} \cdot e^{-t}) = \Delta I(e^{-s}) + \Delta I(e^{-t}) = g(s) + g(t),$

so g satisfies the additive Cauchy equation on $[0, \infty)$. By (L2), g is continuous; by (L3), g is non-decreasing (since e^{-t} is decreasing in t and ΔI is non-increasing in its argument). The unique continuous monotone solution to Cauchy's additive equation is

$g(t) = \lambda t, \lambda \geq 0.$

The case $\lambda = 0$ gives $\Delta I \equiv 0$, which is excluded as a non-trivial measure of information. Hence $\lambda > 0$, and

$\Delta I(x) = g(-\ln x) = -\lambda \ln x$. ■

Justification of the premises

(L1) is the natural extensivity of an information cost: composing two independent extensions multiplies amplitudes (by Proposition 1 and (A4)) and so must add information. (L2) is inherited from continuity of the unitary action on amplitudes. (L3) is the requirement that smaller-amplitude candidates carry larger extension cost, with a full-amplitude candidate ($x = 1$) costing nothing. None of these is a primitive: each follows from prior VERSF structure.

Corollary 1 (The Born rule as the minimal-information extension measure)

With $\lambda = 1$ (fixed by Theorem 1, see below), the extension information of candidate C_i is

$\mathbf{\Delta I_i \equiv -\ln |c_i|^2}$,

and the Born weighting takes the Gibbs form

$\mathbf{p_i = \exp(-\Delta I_i)}$.

It is the unique admissible probability measure on $\mathcal{C}$ whose log-density coincides with the information cost of extending ρ by C_i.

Proof. By Lemma 1, $\Delta I(x) = -\lambda \ln x$. Substituting into $p_i = \exp(-\Delta I_i)$ gives $p_i = |c_i|^{2\lambda}$. Theorem 1 requires the exponent to be 1, fixing $\lambda = 1$. Hence $\Delta I_i = -\ln |c_i|^2$ and $p_i = \exp(-\Delta I_i) = |c_i|^2$. ■

Significance. Corollary 1 unifies four levels of VERSF structure in one identity:

Level	Object	Role in $p_i = \exp(-\Delta I_i)$
Probabilistic	p_i	Born weighting
Informational	ΔI_i	Extension cost on ρ
Geometric	$\partial \Delta I_i / \partial \rho$	Information metric on extension space
Dynamical	κ-field	Sources the gradient of ΔI_i across ρ

Probability, entropy, geometry, and the κ-field are therefore not four independent strands of the framework but four readings of a single Gibbs measure on the admissible extension space of ρ. The detailed coupling to κ-field dynamics is treated in §5.3 and in the companion papers on κ-field structure and on commitment-capacity density.

---

5. Interpretation

5.1 Probability as record consistency

Theorem 1 identifies probability with one physical content: *the relative degree to which admissible observables remain invariant under an admissible record extension C_i, evaluated within reversible pre-commitment structure.* It is not a measure of ignorance, not an objective frequency, and not a degree of belief. It is a structural invariance property of the record's admissible extension space — a geometric property of ρ, read against its possible futures.

5.2 Measurement as commitment selection

In this picture, "wavefunction collapse" is the addition of one C_i to ρ. Probability does not govern *whether* the record extends — that is forced by admissibility, since the universe must produce facts to remain physical — but *which* admissible extension is realised. The weighting $p_i = |c_i|^2$ is the unique measure compatible with the structure of the record. The act of selection is the commitment event itself; nothing else is needed to instantiate it.

5.3 Entropic, geometric, and κ-field readings

Corollary 1 establishes that the Born weighting is itself a Gibbs measure over extension information. This identification has three immediate physical consequences worth stating explicitly.

Entropic. Probability and entropy are not separate quantities related by a postulate; the Born measure *is* the exponential of the negative extension information. This recovers, and gives formal status to, the Born-rule-as-entropic-unfolding result developed elsewhere in the programme.

Geometric. ΔI_i defines a natural distance functional on the extension space $\mathcal{C}$. The associated Fisher-type metric on probability measures over $\mathcal{C}$ is the information geometry within which pre-commitment dynamics flow. Hilbert-space geometry is, on this reading, the information geometry of admissible record extensions.

κ-field. Since the κ-field governs gradients of the commitment-capacity density and ΔI is a functional of ρ, κ-field dynamics induce gradients in the extension information ΔI across the record. Probability is therefore not a static labelling of outcomes but is *dynamically realised* as the flow induced by κ across the admissible extension space. This supplies the formal bridge between the κ-field papers and the present probabilistic structure: probability, entropy, and κ-field dynamics share a single mechanism — gradient flow on $\Delta I[\rho]$ — across the admissible extension space. The κ-field is, in this sense, the dynamical face of the same Gibbs measure that Theorem 1 derives statically.

The two routes to the Born rule developed in the programme — consistency weighting (this paper) and entropic unfolding (companion paper) — are therefore not parallel arguments for the same conclusion but two faces of one identity, made formal by Corollary 1.

5.4 Compatibility with the single-source theorem

Probability is not a new degree of freedom. Theorem 1 shows that the weighting is determined entirely by the relation between the candidate extension and the record under admissibility; Theorem 2 shows that no alternative weighting can exist. Hence

$p_i = \mathcal{F}\_p[\rho]$,

and the single-source identification $\mathcal{O} = \mathcal{F}[\rho]$ holds for probabilistic observables, completing the chain previously established for kinematics, dynamics, geometry, and the fundamental constants.

---

6. Relation to existing derivations

Theorem 1 is structurally a Gleason-type result: a uniqueness theorem for measures on a projection lattice under invariance and additivity assumptions. Four points distinguish it from the standard literature.

(i) The invariance assumptions are derived, not posited. (A2)–(A4) are each consequences of prior VERSF results — single-source structure (giving (A2)), the unitary character of pre-commitment dynamics (giving (A3)), and disjoint-sector decomposition of ρ (giving (A4) via Proposition 1). Each has independent physical content; none is selected for the convenience of the proof. In Gleason's theorem, by contrast, the non-contextuality of the frame functional is introduced as a primitive constraint with no deeper grounding than its mathematical convenience.

(ii) The argument applies in dim $\mathcal{H} = 2$. Gleason's theorem fails in two dimensions; this argument does not, because compositional multiplicativity (A4) is sensitive to the disjoint-sector structure of ρ rather than to the projection lattice of a single Hilbert space. The qubit case is the most physically important — it is the regime of every elementary measurement — and the present argument covers it without modification.

(iii) The probability that emerges has a physical reading. "Record consistency" is a substantive identification, not a formal one. The decision-theoretic derivations of Deutsch and Wallace impose rationality axioms on agents; here no agents appear. Zurek's envariance argument uses environment-induced symmetries; here the relevant symmetry is the unitary structure of pre-commitment dynamics, derived elsewhere in the programme rather than assumed. Frequentist derivations require the hypothesis of repeated trials and an associated typicality measure; here a single commitment event suffices.

(iv) The result is sharpened to a no-go statement. Theorem 2 closes the argument: it is not merely that Born is *one* admissible probability rule, it is that *no other admissible rule can exist*. None of Gleason, Deutsch–Wallace, or Zurek delivers this stronger conclusion, because none rests on a single-source theorem strong enough to forbid alternative probability structures from drawing on data outside the principal physical state.

In each case the move is the same: VERSF replaces an unmotivated postulate or external assumption with a structural feature of the record.

---

7. Scope and limitations

Theorem 1 establishes the Born rule for outcomes of a single commitment event in a finite-dimensional pre-commitment Hilbert space. Several extensions remain open.

- **Continuous spectra.** The argument assumes a discrete orthonormal resolution. Extension to continuous projection-valued measures should follow from the standard limiting procedure but has not been worked out within VERSF.
- **Non-equilibrium weighting.** The theorem assumes that the existing record ρ is in a regime where reversible pre-commitment dynamics dominate prior to commitment — exactly the "preserving reversible structure" clause of the consistency definition. Far-from-equilibrium regimes may modify W through κ-field coupling; this is treated separately.
- **Gauge and interaction sectors.** Compositional closure (Proposition 1, (A4)) is applied here to causally disjoint factor systems. Its action on gauge-coupled subsystems — in particular within the SU(3) × SU(2) × U(1) structure derived from K = 7 hexagonal closure — requires the explicit treatment given in the *Standard Model from Hexagonal Geometry* papers.
- **Multi-time histories.** The present argument is single-event. Iterated commitment histories and their associated path-weighting structure are addressed in the companion work on PAR/CC.

These are extensions, not weaknesses. The result above is sharp within its stated scope.

---

8. Conclusion

Within VERSF, probability is a derived quantity. The Born rule is the unique consistency measure over admissible extensions of the committed record compatible with admissibility, unitary pre-commitment dynamics, and disjoint-sector decomposition (Theorem 1). It is also the *only* admissible probability structure: no alternative is permitted by the single-source theorem (Theorem 2). Its information-theoretic form is forced by the unique extensive measure of extension cost (Lemma 1, Corollary 1). No additional probabilistic axiom is required at any stage.

Probability, in this picture, is the geometry of admissible futures relative to the committed past. It carries no degrees of freedom of its own. It is exactly what the structure of the record permits — and nothing else is permitted at all.

Companion papers within the VERSF programme

- *Why a Fact-Producing Universe Must Satisfy Interference* — derivation of complex Hilbert space from admissibility axioms.
- *The Born Rule as Entropic Unfolding and the Double Square Rule* — independent informational route to $p_i = |c_i|^2$.
- *Pre-Factual Algebraic Reversibility and Compositional Completeness (PAR/CC)* — the algebra underwriting (A4) and Proposition 1.
- *Internal Admissible Closure (IAC)* — the closure conditions underwriting (A2).
- *Commitment-Capacity Density* — the structure of ρ and the No Multi-Primitive Occupancy theorem.
- *Facts as Structural Necessities* — the Joint Necessity Theorem requiring irreversible facts.
- *VERSF Constraint and Lagrangian (BCB)* — the action-principle face of the framework.