The Born Rule from Finite Observation

A Conditional Derivation of the Fixed-Context Born Form

Aernoud Dekker

June 2026; extended July 2026

DOI: 10.17605/OSF.IO/U5RDE

Abstract

This paper gives a conditional derivation of the binary Born probability form $p(\theta) = \cos^2(\theta/2)$ in the physical basis coordinate θ. The derivation has two load-bearing bridge assumptions: that the effective tracking capacity C_{eff} is capacity for preserving operational distinguishability of finite observer records (Fisher capacity bridge), and that the scalar self-ignorance rate $\kappa = h_{\text{KS}} - C_{\text{eff}} \ln 2$ is homogeneous in the physical basis coordinate (scalar-threshold homogeneity). Under the Fisher capacity bridge plus standard finite-resolution invariance, Cencov's theorem selects Fisher–Rao as the unique distinguishability metric on observer records; the binary record then carries the Fisher-arclength identity $p(s) = \cos^2(s/2)$ by elementary geometry. Scalar-threshold homogeneity then forces equal increments of θ to carry equal Fisher cost, so $I(\theta) = \alpha^2$ is constant and the Fisher arclength coordinate is affine in the physical basis coordinate, $s = \alpha\theta$. Endpoint calibrations $p(0) = 1$ and $p(\pi) = 0$, on the first monotone calibration interval, fix $\alpha = 1$ and give $p(\theta) = \cos^2(\theta/2)$. The result extends to the full outcome distribution of a single calibrated n-outcome context: a curvature obstruction shows the homogeneity premise cannot hold jointly for several dials—the record geometry is spherical, and a jointly homogeneous chart would be flat—so no order-free uniform calibration exists, and calibration proceeds instead by hierarchical binary discrimination. Under the conditional form of the same premises, applied per discrimination, the outcome weights are squared components of a real unit vector, $p_i = q_i^2$, for every calibration in the hierarchical dial-local class, with a falsifiable signature: the Fisher information matrix in the calibrated dials is diagonal, each entry equal to the probability mass reaching its discrimination (the node-mass law; for the chain calibration, the unresolved residual mass). The result is conditional on the bridge assumptions. Neither the bridge nor the full Hilbert-space formalism of quantum mechanics is derived here—and a companion limit theorem [1] shows the latter cannot be obtained by the same route, the Markov invariance used here admitting no invariant phase structure.

1 Purpose and IOF Starting Point

The target of this paper is the narrow binary Born form

$$p(\theta) = \cos^2(\theta/2),$$

in the physical basis angle θ of a laboratory measurement, derived from finite observer record constraints together with one substantive physical bridge. The paper does not derive complex Hilbert space, projective composition, tensor products, unitary dynamics, the Standard Model Lagrangian, or the ontic admissible-history measure of the Ignorant Observer Framework (IOF)

[2]. It isolates the binary probability form available once the observer's finite record geometry is fixed and the physical basis coordinate is connected to it; Section 7 then extends the isolation to the full outcome distribution of a single calibrated n-outcome context, while the composition of incompatible contexts remains underived (see the scope discussion below).

IOF [2] treats the measurement basis as a physical reference variable $\theta(t)$ that must be tracked through a finite-rate channel; the bandwidth-limited control realization and its experimental test are developed in the companion BLQC paper [3]. In an interferometer, a spin experiment, or a qubit platform, the setting θ is implemented by a phase reference, pulse phase, local oscillator, magnetic-field direction, or equivalent physical controller. The keystone quantity is the *self-ignorance rate* κ, the gap between the information-production rate of the reference dynamics and the useful tracking capacity:

$$\kappa = h_{\mathrm{KS}} - C_{\mathrm{eff}} \ln 2.$$

Here h_{KS} is an entropy-rate or instability proxy for the reference dynamics, and C_{eff} is the effective channel rate that constrains the reference. The sign of κ is the tracking threshold: in the capacity-wins regime $\kappa < 0$, the observer/controller can in principle keep the reference resolved, and textbook fixed-basis quantum mechanics is operationally available; in the chaos-wins regime $\kappa > 0$, standard quantum mechanics still holds, but the observer works with an unstable, unresolved reference frame whose basis uncertainty grows. IOF by itself describes the resulting visibility attenuation: an observer-relative phase-averaging effect, classical in origin and recoverable in principle when the missing reference information is supplied—not a deviation from standard quantum mechanics. It does not, by itself, explain why the underlying binary probability law is Born. That is the gap addressed here.

The argument proceeds in five steps:

$$\begin{gathered}
\text{IOF finite-rate tracking of physical basis reference } \theta \\
\Downarrow \\
\text{tracking capacity must preserve operational distinguishability of records} \\
\Downarrow \\
\text{Cencov: Fisher–Rao is the unique invariant metric on finite record distributions} \\
\Downarrow \\
\text{scalar threshold } \kappa \text{ requires constant Fisher cost per unit } \theta \\
\Downarrow \\
p(\theta) = \cos^2(\theta/2) \text{ for calibrated binary records.}
\end{gathered}$$

Only one of these steps is a new physical premise: the Fisher capacity bridge (step two). The remainder is either standard information geometry (step three), elementary calculation (steps four and five), or operational structure (step one). The paper states the bridge explicitly, develops the geometry carefully, and shows where the conditional weight lies.

2 Finite Records and Operational Distinguishability

Let a finite observer's accessible record be a probability vector

$$p = (p_1, \ldots, p_n), \qquad p_i \geq 0, \qquad \sum_{i=1}^{n} p_i = 1.$$

This is an operational object. It is not yet a quantum state. It represents the observer's finite distribution over distinguishable record outcomes.

For a BLQC platform with physical basis coordinate θ, the record family is

$$\theta \mapsto p(\theta), \qquad p_o(\theta) = p(o \mid \theta).$$

If a tracking error $d\theta$ changes the record distribution by dp_o, the operational size of that error is not the coordinate displacement alone. It is the distinguishability of the changed record. Two reference errors of the same coordinate magnitude but with different effects on $p(o \mid \theta)$ are not operationally equal for the purposes of finite-rate tracking. This observation is the entry point of the IOF–Fisher bridge.

2.1 The Fisher Capacity Bridge

Assumption 2.1 (Fisher capacity bridge)**.** The effective tracking capacity C_{eff} measures the useful rate of reducing distinguishability error in the operational record family $p(o \mid \theta)$. Therefore the tracking error relevant to IOF is measured in the invariant distinguishability metric of finite record distributions.

This is the first and most substantive physical premise of the paper. It says that useful tracking bits are not abstract bits detached from the experiment. They are bits that reduce uncertainty in the operational records that define the basis. A reference error matters only insofar as it changes the observer's probability records. Therefore the natural metric for the tracking task is not an arbitrary Euclidean metric on the control dial, but the invariant distinguishability metric on finite record distributions. The next two sections show that this metric is uniquely Fisher–Rao.

2.2 Finite Resolution and Markov Kernels

Finite capacity implies that records can be coarse-grained or refined. At the level of probability distributions, an ordinary loss of finite resolution is represented by a Markov kernel

$$p'(m') = \sum_m K(m' \mid m)\, p(m),$$

where

$$K(m' \mid m) \geq 0, \qquad \sum_{m'} K(m' \mid m) = 1.$$

Definition 2.2 (Nested finite-resolution reduction)**.** Let $R_C : A \to M_C$ be a record projection from an admissible history space A to a finite record alphabet M_C. A lower-capacity projection $R_{C'} : A \to M_{C'}$ is *nested* in R_C when there is a deterministic map $q_{C'|C} : M_C \to M_{C'}$ such that

$$R_{C'} = q_{C'|C} \circ R_C.$$

Proposition 2.3 (Nested reductions induce Markov kernels)**.** *If $R_{C'} = q_{C'|C} \circ R_C$, then the distribution of lower-resolution records is obtained from the distribution of higher-resolution records by a Markov kernel.*

Proof. Let μ_θ be a measure over admissible histories at basis setting θ. The high-resolution record distribution is the pushforward

$$p_C(m \mid \theta) = \mu_\theta\{a \in A : R_C(a) = m\}.$$

The lower-resolution distribution is

$$p_{C'}(m' \mid \theta) = \mu_\theta\{a \in A : R_{C'}(a) = m'\}.$$

Using the nesting relation,

$$\{a : R_{C'}(a) = m'\} = \bigcup_{m:\, q_{C'|C}(m)=m'} \{a : R_C(a) = m\},$$

where the union is disjoint. Therefore

$$p_{C'}(m' \mid \theta) = \sum_m K_{C'|C}(m' \mid m)\, p_C(m \mid \theta),$$

with

$$K_{C'|C}(m' \mid m) = \begin{cases} 1, & q_{C'|C}(m) = m', \\ 0, & q_{C'|C}(m) \neq m'. \end{cases}$$

This kernel is stochastic, so the induced finite-resolution reduction is Markov. □

Refinement is the reverse operation. A refinement is *sufficient* when it changes the record alphabet without adding new information about θ. For example,

$$p_i \mapsto (ap_i, (1-a)p_i),$$

with fixed $a \in [0, 1]$ independent of θ, is sufficient because the original record is recovered by summing the two refined components, and the split itself carries no θ-information.

Assumption 2.4 (Finite-resolution invariance)**.** Admissible finite-resolution changes of observer records are Markov maps, and sufficient refinements or coarse-grainings preserve distinguishability of the tracked basis family $\{p(o \mid \theta)\}$.

Remark 2.5 (Sufficiency versus lossy coarsening)**.** The invariance premise is not the claim that every finite-capacity loss preserves distinguishability. Generic coarse-graining can destroy information about θ, and such maps are not symmetries of the record geometry. The invariance is invoked only for resolution changes that alter the representation of the record without losing the tracked parameter. In statistics, those are precisely sufficient Markov maps. Lossy capacity degradation is a separate physical process; it belongs to the IOF coarsening dynamics and to the empirical visibility-loss protocol.

3 Cencov's Theorem and Fisher–Rao Geometry

Let

$$\Delta_n^\circ = \left\{ p \in \mathbb{R}^n : p_i > 0,\ \sum_i p_i = 1 \right\}$$

be the interior of the finite probability simplex. Its tangent vectors $u \in T_p\Delta_n^\circ$ satisfy $\sum_i u_i = 0$. A Markov kernel K maps distributions and tangent vectors by

$$(Kp)_a = \sum_i K(a \mid i)p_i, \qquad (Ku)_a = \sum_i K(a \mid i)u_i.$$

A Markov map is *sufficient* for a statistical model $\{p_\theta\}$ when there is a recovery Markov map L such that $LKp_\theta = p_\theta$ for all relevant θ. Differentiating along the model gives the corresponding recovery of tangent directions.

Theorem 3.1 (Cencov uniqueness, finite form). *Let $g^{(n)}$ be a smooth family of Riemannian metrics on the interiors Δ_n°. If this family is invariant under sufficient Markov morphisms, in the sense that*

$$g_{Kp}^{(m)}(Ku, Kv) = g_p^{(n)}(u, v)$$

whenever K is sufficient for the model directions under consideration, then $g^{(n)}$ is, up to an overall positive constant, the Fisher–Rao metric:

$$g_p(u, v) = c \sum_i \frac{u_i v_i}{p_i}.$$

A proof is given by Cencov [4]; modern treatments appear in Amari and Nagaoka [5].

Combining Theorem 3.1 with Assumption 2.4 (finite-resolution invariance) and Assumption 2.1 (Fisher capacity bridge), the observer's record manifold carries the Fisher–Rao metric

$$ds^2_{\mathrm{FR}} = c \sum_i \frac{dp_i^2}{p_i}.$$

The constant $c > 0$ fixes units; in what follows it is set to 1.

The important point is directional. Fisher–Rao geometry is not selected because it resembles quantum mechanics. It is selected because it is the unique distinguishability geometry compatible with sufficient finite-resolution changes of classical records. Quantum structure has not yet been assumed. In IOF vocabulary: the metric in which the basis-tracking task is correctly measured is Fisher–Rao, because the useful component of C_{eff} is, by Assumption 2.1, capacity for reducing Fisher distinguishability error.

4 Square-Root Coordinates and the Binary Fisher Identity

Define square-root record coordinates

$$q_i = \sqrt{p_i}.$$

Since $\sum_i p_i = 1$, these coordinates satisfy

$$\sum_i q_i^2 = 1,$$

so $q = (q_1, \ldots, q_n)$ lies on the positive orthant of the unit sphere. Differentiating $q_i = \sqrt{p_i}$ gives

$$dq_i = \frac{dp_i}{2\sqrt{p_i}},$$

and therefore

$$4 \sum_i dq_i^2 = \sum_i \frac{dp_i^2}{p_i} = ds^2_{\mathrm{FR}}.$$

Thus Fisher–Rao record geometry is spherical geometry in square-root coordinates, up to the conventional factor of 4. The square root is not inserted as a quantum amplitude. It is the coordinate induced by the invariant statistical metric.

For a binary record,

$$p = (p, 1 - p), \qquad 0 < p < 1,$$

the Fisher–Rao line element specialises to

$$ds^2_{\mathrm{FR}} = \frac{dp^2}{p} + \frac{d(1-p)^2}{1-p} = \frac{dp^2}{p(1-p)}.$$

Introduce the Fisher arclength coordinate

$$s = 2 \arccos \sqrt{p}.$$

Then

$$p = \cos^2\left(\frac{s}{2}\right), \qquad 1-p = \sin^2\left(\frac{s}{2}\right).$$

Direct differentiation gives $dp = -\frac{1}{2}\sin s\, ds$, while $p(1-p) = \frac{1}{4}\sin^2 s$. Therefore

$$\frac{dp^2}{p(1-p)} = ds^2,$$

confirming that s is an arclength coordinate of the Fisher–Rao metric on the binary simplex.

Proposition 4.1 (Binary Fisher geometry)**.** *For a calibrated binary record with endpoints $p = 1$ and $p = 0$, the Fisher arclength coordinate $s = 2\arccos\sqrt{p}$ satisfies*

$$p(s) = \cos^2\left(\frac{s}{2}\right), \qquad 1 - p(s) = \sin^2\left(\frac{s}{2}\right),$$

with $s \in [0, \pi]$.

Proof. The coordinate $s = 2\arccos\sqrt{p}$ is an arclength coordinate on $0 < p < 1$ because $ds^2_{\mathrm{FR}} = ds^2$. The endpoints are limiting calibrations: $p \to 1$ gives $s \to 0$, and $p \to 0$ gives $s \to \pi$. Inversion gives the stated result. □

This is the binary Born form *in the Fisher arclength coordinate s*. It is geometry, not dynamics: once s is defined as Fisher–Rao arclength, the identity $p(s) = \cos^2(s/2)$ is an elementary consequence of binary record geometry under the Fisher–Rao metric. What remains is to connect s to the physical basis coordinate θ of the BLQC experiment. This is the work of the next section.

5 The Scalar Threshold κ Forces Constant Fisher Information in θ

The Fisher arclength coordinate s is a coordinate on the record manifold. The physical basis coordinate θ is a coordinate on the laboratory control dial. These are different coordinates on different spaces. The map $\theta \mapsto p(\theta)$ relates them, but does not by itself fix any particular relation $s(\theta)$. For the IOF–Fisher bridge to close, a further physical premise about how θ and s are related is required. The premise comes from the scalar form of the self-ignorance rate κ itself.

The self-ignorance rate κ is written as a scalar:

$$\kappa = h_{\mathrm{KS}} - C_{\mathrm{eff}} \ln 2.$$

It is *not* written as a position-dependent quantity

$$\kappa(\theta) = h_{\mathrm{KS}}(\theta) - C_{\mathrm{eff}}(\theta) \ln 2,$$

and it is not written with a position-dependent conversion between physical basis increments and record distinguishability. The scalar form is a structural feature of the IOF framework: it claims that one threshold controls the tracking task across the calibrated basis range.

By Assumption 2.1, useful tracking capacity has been identified with Fisher distinguishability capacity. If the scalar threshold κ is to be the right characterisation of the tracking task across the basis range, then equal increments of the physical basis coordinate θ must impose equal Fisher distinguishability cost. Otherwise the threshold would be effectively local in θ, and the scalar form would be a coordinate-dependent artefact rather than a structural claim.

The scalar form of κ is not a mathematical identity; it is an empirical hypothesis about the physical reference channel. A position-dependent threshold $\kappa(\theta)$ would be logically possible and would falsify the homogeneity assumption in the original basis coordinate.

Equivalently, scalar-threshold homogeneity is a calibration consistency condition. If one scalar threshold is claimed to characterize basis tracking across the calibrated operational range, then the chosen physical basis coordinate must already be a coordinate in which equal increments impose equal operational tracking burden. Under the Fisher capacity bridge, operational tracking burden is Fisher distinguishability burden. Hence the coordinate must be Fisher-affine.

Assumption 5.1 (Scalar-threshold homogeneity)**.** For a calibrated basis reference, the scalar threshold $\kappa = h_{\text{KS}} - C_{\text{eff}} \ln 2$ requires that the physical basis coordinate θ is homogeneous with respect to the Fisher distinguishability metric on observer records.

The consequence is direct. Homogeneity in θ of the Fisher cost means

$$ds = \alpha\, d\theta$$

for some constant $\alpha > 0$, independent of θ. Equivalently, in terms of Fisher information,

$$I(\theta) = \left(\frac{ds}{d\theta}\right)^2 = \alpha^2.$$

Remark 5.2 (The substantive new premise)**.** This is the load-bearing assumption of the bridge. Without it, IOF can still predict basis-uncertainty visibility loss, but it does not force the Born probability law. A non-Born record law would generally make $I(\theta)$ position-dependent, requiring either a local $\kappa(\theta)$ or a non-homogeneous basis coordinate. Reading the self-ignorance rate κ as genuinely scalar in a calibrated basis is, therefore, the bridge premise. Both Assumption 2.1 and Assumption 5.1 are explicitly conditional: they are stated, used, and exposed as the point where the derivation could fail empirically. See §9.

6 The Binary Born Form in θ

The two sides of the bridge are now in place. Proposition 4.1 gives the binary record law in the Fisher arclength coordinate s:

$$p(s) = \cos^2(s/2).$$

Assumption 5.1 gives the relation between s and the physical basis coordinate θ:

$$s = \alpha\theta + s_0$$

for constants $\alpha > 0$ and s_0, obtained by integrating $ds = \alpha\, d\theta$. Substituting,

$$p(\theta) = \cos^2\left(\frac{\alpha\theta + s_0}{2}\right).$$

The calibration $p(0) = 1$ gives $s_0 = 0$, so

$$p(\theta) = \cos^2\left(\frac{\alpha\theta}{2}\right).$$

The opposite endpoint calibration $p(\pi) = 0$ requires $\cos(\alpha\pi/2) = 0$, so α is any positive odd integer; choosing the first monotone calibration interval fixes $\alpha = 1$. Therefore

$$p(\theta) = \cos^2(\theta/2).$$

For completeness, the same conclusion can be derived directly from $I(\theta) = \alpha^2$ and the binary Fisher identity. For a binary record,

$$I(\theta) = \frac{(dp/d\theta)^2}{p(1-p)}.$$

Setting $I(\theta) = \alpha^2$ and choosing the monotone branch from certainty to impossibility,

$$\frac{dp}{d\theta} = -\alpha\sqrt{p(1-p)},$$

which separates as

$$\frac{dp}{\sqrt{p(1-p)}} = -\alpha\, d\theta.$$

Since $\int dp/\sqrt{p(1-p)} = 2\arcsin\sqrt{p}$, integration gives

$$2\arcsin\sqrt{p(\theta)} = -\alpha\theta + C.$$

With $p(0) = 1$ we have $C = \pi$, hence $\sqrt{p(\theta)} = \cos(\alpha\theta/2)$. The calibration $p(\pi) = 0$ again requires $\cos(\alpha\pi/2) = 0$; choosing the first monotone calibration interval fixes $\alpha = 1$, and

$$p(\theta) = \cos^2(\theta/2).$$

The two derivations are equivalent: the first uses the geometric identity of §4 and the bridge map of §5; the second integrates the constant-Fisher-information ODE directly. Both expose where the conditional weight lies.

Theorem 6.1 (Conditional binary Born form under IOF)**.** *Suppose that:*

1. *the effective tracking capacity C_{eff} is capacity for preserving operational distinguishability of finite observer records (Assumption 2.1);*
2. *admissible finite-resolution changes are sufficient Markov maps and preserve distinguishability of the tracked basis family (Assumption 2.4);*
3. *the calibrated basis reference has a scalar homogeneous threshold $\kappa = h_{\text{KS}} - C_{\text{eff}} \ln 2$ in the physical basis coordinate θ (Assumption 5.1);*
4. *the calibration map $\theta \mapsto p(\theta)$ is monotone on $[0, \pi]$, with endpoints $p(0) = 1$ (certainty) and $p(\pi) = 0$ (impossibility).*

Then the calibrated binary record carries the Born form

$$p(\theta) = \cos^2(\theta/2).$$

Proof. Assumptions 2.1 and 2.4 with Cencov's theorem (Theorem 3.1) place the Fisher–Rao metric on the observer's record manifold. The binary specialisation (Proposition 4.1) gives $p(s) = \cos^2(s/2)$ in Fisher arclength. Assumption 5.1 gives $ds = \alpha\, d\theta$, so $s = \alpha\theta + s_0$. The calibrations of clause (4) fix $s_0 = 0$ and $\alpha = 1$, yielding $p(\theta) = \cos^2(\theta/2)$. □

7 The Multi-Outcome Extension: One Fixed Context

The binary result raises an immediate question: does the same route reach the outcome distribution of an n-outcome measurement context? The answer has two parts, and both carry weight. First, the binary homogeneity premise does *not* generalize jointly: for $n \geq 3$ there is no calibration in which several dials are everywhere uniform at once, and the failure is a curvature obstruction, not a technical gap. Second, the premise that survives—conditional calibration, one binary discrimination at a time, organized on any hierarchy—forces the outcome weights to be squared components of a real unit vector, for the full outcome distribution of the context and for every calibration topology in the hierarchical dial-local class. Everything in this section concerns a *single* calibrated context. The consistency of weights across incompatible contexts is not derived here, and the companion limit theorem [1] shows that the phase structure governing context composition cannot be generated by this route at all.

7.1 The obstruction

A *calibration chart* for an n-outcome context is a C^2 map $\Phi : U \to \Delta_n^\circ$, with $U \subseteq \mathbb{R}^{n-1}$ open, from physical dial coordinates $\theta = (\theta^1, \dots, \theta^{n-1})$ to record distributions. The joint generalization of Assumption 5.1 would demand that equal increments of every dial impose equal Fisher cost in every direction: the pullback $\Phi^* g_{\mathrm{FR}}$ would be the Euclidean metric $\alpha^2 \delta_{ab}$ on U.

Theorem 7.1 (No jointly homogeneous multi-dial calibration). *For $n \geq 3$ there is no calibration chart whose pullback of the Fisher–Rao metric is Euclidean.*

Proof. In square-root coordinates, $(\Delta_n^\circ, g_{\mathrm{FR}})$ is isometric to an open subset of the radius-2 round sphere (§4), hence has constant sectional curvature $1/4 > 0$. Sectional curvature is invariant under Riemannian isometry, and the Euclidean metric is flat. □

Remark 7.2 (Why the binary case is exceptional). For $n = 2$ the record manifold is one-dimensional and carries no intrinsic curvature; the scalar premise of Assumption 5.1 is exactly the flat chart that exists only in dimension one. The binary case was not the fortunate small case; it was the only case in which joint homogeneity is even consistent.

Remark 7.3 (Physical reading). A finite observer cannot spread an n-way discrimination over $n-1$ dials that are all uniform at once. Three claims of different strengths must be kept apart here. What the theorem forces is only the failure of order-free uniformity: no calibration treats all dials alike everywhere. What is then *motivated*, not forced, is the class of hierarchical binary calibrations—decompositions of the n-way discrimination into binary discriminations, the structure a C_{eff}-limited channel realizes a bounded number of at a time. And what is again a theorem (Corollary 7.11 below) is that within that class the derived form does not depend on which hierarchy is chosen. The particular chain used below is a canonical representative, a protocol choice; the form is not.

7.2 Nested calibration

Definition 7.4 (Nested calibration). A *nested calibration* of an n-outcome context is an ordering of the outcomes $1, \dots, n$ together with dials $\theta^1, \dots, \theta^{n-1}$ such that dial θ^k calibrates the two-outcome *conditional* record $\{k\}$ versus $\{k+1, \dots, n\}$, conditional on the record event "not among outcomes $1, \dots, k-1$". Writing $c_k := p(k \mid \text{not} < k)$ for the conditional weights, the calibration is *dial-local*: $\partial c_j / \partial \theta^k = 0$ for $j \neq k$.

Dial-locality is the operational content of "calibrates the conditional record", and it is realizable by ordinary laboratory pulse sequences (Remark 7.13 below).

Definition 7.5 (Hierarchical dial-local calibration)**.** More generally, let T be a full binary tree whose leaves are the outcomes $1, \dots, n$, with a dial θ^ν attached to each internal node ν and c_ν the conditional weight of ν's left subtree given that the outcome reaches ν. A *hierarchical dial-local calibration* on T is one in which each dial moves only its own node's conditional weight. The nested calibration of Definition 7.4 is the chain (caterpillar-tree) instance.

Remark 7.6 (The premise change from the binary case)**.** In the binary case, scalar-threshold homogeneity is a single global claim: one threshold across the calibrated range. In the extension it is applied *per discrimination*: each internal node is its own binary reference-tracking task, with its own scalar threshold, defined conditionally on the discriminations above it. That is a weaker premise than a global one, and it is the honest price of the extension—Theorem 7.1 shows that no global homogeneity claim is available to take its place. Lemma 7.9(iv) shows the per-node claim is well posed (the two natural readings coincide along each dial), and the module of §9.1 exposes it to measurement directly.

7.3 The conditional decomposition of the record metric

Lemma 7.7 (Conditional product structure)**.** *In conditional coordinates $c = (c_1, \dots, c_{n-1})$, the Fisher–Rao metric on Δ_n° is the orthogonal sum*

$$ds^2_{\mathrm{FR}} = \sum_{k=1}^{n-1} r_k \, \frac{dc_k^2}{c_k(1-c_k)}, \qquad r_k = \prod_{j<k} (1 - c_j),$$

where each factor $dc^2/(c(1-c))$ is the binary Fisher–Rao metric of §4, and $r_k = 1 - \sum_{j<k} p_j$ is the unresolved residual mass before layer k.

Proof. With $p_k = c_k r_k$ for $k \le n-1$ and $p_n = r_n$, the record's log-likelihood factorizes sequentially,

$$\ln p_X = \sum_k \Big[\mathbf{1}\{X = k\} \ln c_k + \mathbf{1}\{X > k\} \ln(1 - c_k) \Big],$$

with scores $\partial_{c_k} \ln p_X = \mathbf{1}\{X = k\}/c_k - \mathbf{1}\{X > k\}/(1 - c_k)$ and moments $\mathbb{E}\,\mathbf{1}\{X = k\} = c_k r_k$, $\mathbb{E}\,\mathbf{1}\{X > k\} = (1 - c_k) r_k$. For $k < l$, using $\mathbf{1}\{X = k\}\mathbf{1}\{X = l\} = 0$, $\mathbf{1}\{X = k\}\mathbf{1}\{X > l\} = 0$, $\mathbf{1}\{X > k\}\mathbf{1}\{X = l\} = \mathbf{1}\{X = l\}$, and $\mathbf{1}\{X > k\}\mathbf{1}\{X > l\} = \mathbf{1}\{X > l\}$,

$$g_{kl} = -\frac{1}{1 - c_k} \left[\frac{\mathbb{E}\,\mathbf{1}\{X = l\}}{c_l} - \frac{\mathbb{E}\,\mathbf{1}\{X > l\}}{1 - c_l} \right] = -\frac{r_l - r_l}{1 - c_k} = 0,$$

and on the diagonal $g_{kk} = c_k r_k / c_k^2 + (1 - c_k) r_k / (1 - c_k)^2 = r_k / (c_k(1 - c_k))$. ∎

The decomposition has a plain operational reading: each layer's distinguishability budget is the probability mass the earlier discriminations left unresolved. Information about finer distinctions is available only in proportion to what remains undecided. The conditional coordinates themselves are classical: they are the neutrality (stick-breaking) coordinates of the analysis of proportions [6]; the lemma states their Fisher–Rao orthogonality in calibration form.

The decomposition is not special to the chain.

Lemma 7.8 (Node-mass decomposition on trees)**.** *For a hierarchical dial-local calibration on a full binary tree T (Definition 7.5), the Fisher–Rao metric is the orthogonal sum over internal nodes,*

$$ds^2_{\mathrm{FR}} = \sum_{\nu} m_\nu \, \frac{dc_\nu^2}{c_\nu(1-c_\nu)},$$

where m_ν is the node mass*, the total probability of the outcomes below ν. Lemma 7.7 is the chain instance, with $m_k = r_k$.*

Proof. The log-likelihood factorizes along the root-to-leaf path, $\ln p_X = \sum_\nu \big[\mathbf{1}\{X \in L_\nu\} \ln c_\nu + \mathbf{1}\{X \in R_\nu\} \ln(1-c_\nu)\big]$, with L_ν, R_ν the leaf sets of ν's left and right subtrees, scores $\partial_{c_\nu} \ln p_X = \mathbf{1}\{L_\nu\}/c_\nu - \mathbf{1}\{R_\nu\}/(1-c_\nu)$, and moments $\mathbb{E}\,\mathbf{1}\{L_\nu\} = c_\nu m_\nu$, $\mathbb{E}\,\mathbf{1}\{R_\nu\} = (1-c_\nu)m_\nu$. For distinct nodes there are two cases. If neither is an ancestor of the other, the supports are disjoint and every product of indicators vanishes, so the cross term is zero. If ν' lies in, say, the left subtree of ν, then $\mathbf{1}\{L_{\nu'}\} \le \mathbf{1}\{L_\nu\}$ and $\mathbf{1}\{R_{\nu'}\} \le \mathbf{1}\{L_\nu\}$, so

$$g_{\nu\nu'} = \frac{1}{c_\nu}\left[\frac{\mathbb{E}\,\mathbf{1}\{L_{\nu'}\}}{c_{\nu'}} - \frac{\mathbb{E}\,\mathbf{1}\{R_{\nu'}\}}{1-c_{\nu'}}\right] = \frac{m_{\nu'} - m_{\nu'}}{c_\nu} = 0,$$

the same cancellation as in the chain. On the diagonal, $g_{\nu\nu} = m_\nu/(c_\nu(1-c_\nu))$. □

Lemma 7.9 (The layers inherit the record premises)**.** *Under Definition 7.4, each layer k is a two-outcome record family in the sense of §§2–4 for its dial θ^k: (i) merging the outcomes $\{k{+}1,\ldots,n\}$ is a sufficient Markov coarse-graining for the θ^k-family, since by dial-locality the merged block's internal ratios are θ^k-independent; (ii) conditioning on the record event $E_k = \{\text{not} < k\}$ discards no θ^k-information, since $p(E_k) = r_k$ and the record outside E_k are both θ^k-independent; (iii) sufficient Markov recodings of the layer record extend, by the identity on the selector and the other layers, to sufficient Markov recodings of the full record, so Assumption 2.4 applies to the layer; and (iv) by Lemma 7.7 the invariant metric's layer block is the binary Fisher–Rao metric scaled by r_k, which is constant along θ^k, so a per-layer homogeneity premise is well posed: constancy of the conditional Fisher information along the dial and constancy of the full-record Fisher information along the dial are the same condition, the two differing by the θ^k-independent factor r_k.*

Proof. (i) and (ii) are immediate from dial-locality and the product structure: the θ^k-dependence of the full distribution sits entirely in the factor $c_k(\theta^k)$. For (iii), represent the record within E_k as the pair (B_k, Y), where B_k is the layer's binary value ($\{k\}$ versus $\{> k\}$) and Y is the residual detail, defined when B_k takes the value "rest". By dial-locality the joint distribution factorizes as $p(\text{selector}) \cdot p(B_k \mid E_k) \cdot p(Y \mid \text{rest})$ with only the middle factor θ^k-dependent. A Markov recoding K of the layer record, with recovery L, extends to the full record as $\mathrm{id} \otimes K \otimes \mathrm{id}$ (identity on the selector and on Y), with recovery $\mathrm{id} \otimes L \otimes \mathrm{id}$; the extension is Markov, recovers the family, and is sufficient for the θ^k-family precisely because the θ^k-dependence lives in the recoded factor. (iv) restates Lemma 7.7 together with $\partial r_k/\partial\theta^k = 0$. □

7.4 The nested Born form

Theorem 7.10 (Conditional fixed-context Born form)**.** *Suppose that Assumptions 2.1 and 2.4 hold; that the context carries a nested calibration (Definition 7.4); that each conditional dial satisfies scalar-threshold homogeneity (Assumption 5.1 applied to the layer family); and that* *each layer is endpoint-calibrated,* $c_k(0) = 1$ *and* $c_k(\pi) = 0$ *on the first monotone interval. Then the calibrated* n*-outcome record carries*

$$p_k(\theta) = \cos^2\left(\frac{\theta^k}{2}\right) \prod_{j<k} \sin^2\left(\frac{\theta^j}{2}\right) \quad (k \le n-1), \qquad p_n(\theta) = \prod_{j=1}^{n-1} \sin^2\left(\frac{\theta^j}{2}\right),$$

equivalently $p_i(\theta) = q_i(\theta)^2$ *with* $q(\theta)$ *the nested spherical unit vector: the outcome weights are squared components of a real unit vector whose dial dependence is Fisher-arclength-affine at every layer—the fixed-context Born form.*

Proof. By Lemma 7.9, each layer satisfies the hypotheses of Theorem 6.1, so $c_k(\theta^k) = \cos^2(\theta^k/2)$. The chain rule $p_k = c_k \prod_{j<k}(1 - c_j)$ assembles the display, and the identification $p_i = q_i^2$ is the algebra of the nested product. □

Corollary 7.11 (Class invariance)**.** *The same conclusion holds for every hierarchical dial-local calibration (Definition 7.5): under per-node scalar-threshold homogeneity and per-node endpoint calibration,* $c_\nu(\theta^\nu) = \cos^2(\theta^\nu/2)$ *at every internal node, and the leaf weights are squared components of the tree-parametrized real unit vector. The derived form is invariant across calibration topologies; what varies with the topology is the chart and the module law of §9.1.*

Proof. Lemma 7.9 applies verbatim at each internal node, with E_ν the event of reaching ν and the merged blocks the two subtrees; Lemma 7.8 supplies the metric decomposition with the θ^ν-independent weight m_ν. Each node is then an instance of Theorem 6.1, and multiplying the conditional weights down each root-to-leaf path gives $p_i = q_i^2$. □

Remark 7.12 (Class invariance of the form; topology of the chart)**.** Different outcome orderings of the chain give charts related by coordinate permutations—genuine isometries of the ambient sphere. Different trees are simply different hierarchical parametrizations of the same record orthant. In every case the derived form—squared components of a real unit vector—is the same; what is topology-dependent is the chart, and with it the specific node-mass predictions of the module. The topology is fixed by the experimenter's calibration protocol and is therefore known by construction.

Remark 7.13 (Worked realizations)**.** The premise class is nonempty and realized by ordinary quantum hardware. (i) A spin-1 system with $|m{=}0\rangle$ rotated by an angle θ has $p_0 = \cos^2\theta$ and $p_\pm = \frac{1}{2}\sin^2\theta$, with $I(\theta) = 4$ constant and the square-root curve a great circle of the record sphere. (ii) A qutrit prepared by the pulse sequence $U_{23}(\theta^2)\, U_{12}(\theta^1)\, |1\rangle$ has amplitudes exactly the nested spherical vector, with $c_1 = \cos^2(\theta^1/2)$ moved only by the first pulse and $c_2 = \cos^2(\theta^2/2)$ only by the second: the second rotation acts inside the $\{2,3\}$ block and cannot change p_1, while the first fixes only that block's input direction. Sequential pulse stacks are dial-local nested calibrations.

8 Relation to IOF Visibility

IOF's empirical content concerns unresolved basis error—an observer-relative phase-averaging effect within standard quantum mechanics. Under $ds = \alpha\, d\theta$, basis uncertainty can be expressed either as a physical angle uncertainty σ_θ or as a Fisher distinguishability uncertainty σ_s:

$$\sigma_s^2 = \alpha^2 \sigma_\theta^2.$$

In the calibrated normalisation $\alpha = 1$, the IOF visibility growth law

$$V_{\mathrm{IOF}}(t) = \exp\bigl[-\tfrac{1}{2}\sigma_0^2\, e^{2\kappa t}\bigr]$$

—a closure ansatz in the foundational paper [2], not a derived result—can be read as visibility loss from growing uncertainty in the same Fisher-homogeneous coordinate that fixes the binary probability law.

The visibility-modulated binary record is the Born weight read at finite contrast. The link is more than a shared coordinate: the finite-contrast record is the binary Born weight averaged over the unresolved basis. Writing the weight in fringe form,

$$p(\theta) = \cos^2(\theta/2) = \tfrac{1}{2}\bigl(1 + \cos\theta\bigr),$$

and averaging over a realized basis $\theta = \theta_0 + \delta\theta$ with $\delta\theta \sim \mathcal{N}(0, \sigma_\theta^2)$ gives

$$\langle p \rangle = \tfrac{1}{2}\Bigl(1 + e^{-\sigma_\theta^2/2}\cos\theta_0\Bigr) = \tfrac{1}{2}\bigl(1 + V_{\mathrm{IOF}}\cos\theta_0\bigr), \qquad V_{\mathrm{IOF}} = e^{-\sigma_\theta^2/2}.$$

The Gaussian average attenuates only the coherent fringe, by exactly V_{IOF}. The Born weight is therefore the law at full contrast ($\sigma_\theta \to 0$, $V_{\mathrm{IOF}} \to 1$); V_{IOF} is the contrast at which a finite-capacity observer reads that same weight, flattening to the uninformative $\frac{1}{2}$ as $V_{\mathrm{IOF}} \to 0$. This is meaningful only when $\delta\theta$ is genuine unresolved physical variation of the realized reference, not private uncertainty about a sharp value. Ordinary dephasing attenuates the same fringe through an independent angular smearing, so the two convolve and their characteristic functions multiply: the observed visibility is $V_{\mathrm{obs}} = V_{\mathrm{std}}\, V_{\mathrm{IOF}}$. The factorisation is ordinary convolution of independent angular smearings, not a new physical channel; and the two factors differ in kind. V_{IOF} is observer-relative phase averaging, restorable in post-processing whenever a record of the realized reference is supplied—the hallmark of reference-frame physics within quantum mechanics [7]—whereas environmental decoherence is irreversible in the relevant operational sense.

Thus, under the bridge, IOF does not merely add noise to an otherwise unrelated probability rule. The same operational geometry determines:

- the calibrated binary law $p(\theta) = \cos^2(\theta/2)$;
- the visibility factor $V_{\mathrm{IOF}} = e^{-\sigma_\theta^2/2}$, that same law read at finite contrast;
- the metric in which finite-rate basis tracking succeeds or fails;
- the self-ignorance rate $\kappa = h_{\mathrm{KS}} - C_{\mathrm{eff}} \ln 2$.

These are tied to a single object: the Fisher geometry of the observer's binary record family.

9 Empirical Status of the Bridge Assumptions

The bridge assumptions are empirical claims about a physical reference channel, and the BLQC benchmark [3] includes a dedicated *Fisher-homogeneity module* that measures them directly. Operationally, the module measures the record family $p(o \mid \theta)$, computes

$$I(\theta) = \sum_o \frac{1}{p(o \mid \theta)} \left(\frac{\partial p(o \mid \theta)}{\partial \theta}\right)^2,$$

and tests whether $I(\theta)$ is approximately constant over the calibrated reference range (Assumption 5.1).

When the module is run on real hardware, the operative threshold quantity is the calibrated form $\kappa_{\rm cal} = h_{\rm KS} - \eta\, C_{\rm eff} \ln 2$ with the efficiency factor η frozen, deferring to the BLQC benchmark [3] for its calibration.

The protocol also requires reporting unresolved reference error in Fisher units,

$$\sigma_s^2 = I(\theta)\sigma_\theta^2,$$

or, over larger intervals,

$$s(\theta) = \int^{\theta} \sqrt{I(u)}\, du.$$

One point of discipline must be stated plainly. Constant $I(\theta)$ is also a prediction of standard quantum mechanics: for an ideal calibrated binary measurement, $p(\theta) = \cos^2(\theta/2)$ gives $I(\theta) = 1$ identically. The Fisher-homogeneity module is therefore a consistency check, not a discriminator. Confirming constant $I(\theta)$ confirms what quantum mechanics already predicts, and is compatible with the bridge; it cannot single the bridge out against quantum mechanics. The module's empirical force runs in the falsifying direction only: a measured position-dependent $I(\theta)$ in a calibrated record family would reject scalar-threshold homogeneity in that coordinate, and with it the derivation given here.

For the same reason, a κ-scaled visibility breakdown carries no evidential weight for the bridge. The κ control law is classical physics—the data-rate theorem applied to reference tracking—and is fully consistent with standard quantum mechanics, so observing it confirms an operational control law, not a foundational premise. The bridge stands or falls with its two named assumptions, and the only empirical purchase on them is the homogeneity measurement above, in its falsifying direction. The paper is foundations work, not experiment, and its claims should be weighed accordingly.

9.1 The multi-dial module

Theorem 7.10 and Corollary 7.11 upgrade the module's observable from the scalar $I(\theta)$ to the Fisher information *matrix* in the calibrated dials,

$$G_{ab}(\theta) = \sum_o \frac{1}{p(o \mid \theta)} \frac{\partial p(o \mid \theta)}{\partial \theta^a} \frac{\partial p(o \mid \theta)}{\partial \theta^b}.$$

The prediction under the per-node premises is the *node-mass law*, stated relative to the calibration tree the experimenter has built:

$$G_{\nu\nu'}(\theta) = 0 \quad (\nu \neq \nu'), \qquad G_{\nu\nu}(\theta) = m_\nu(\theta),$$

each dial's information equal to the probability mass reaching its discrimination. For the chain calibration this is the *residual-mass law*, $G_{kk} = 1 - \sum_{j<k} p_j$, with the binary module recovered as the top layer, $G_{11} = 1$.

The finite-contrast null generalizes layer-wise. Modelling per-layer imperfection by the fringe form $c_k = (1 + V_k \cos\theta^k)/2$,

$$G_{kk}(\theta) = r_k^V \frac{V_k^2 \sin^2\theta^k}{1 - V_k^2 \cos^2\theta^k}, \qquad r_k^V = \prod_{j<k}(1 - c_j),$$

and the off-diagonal entries remain zero: layer-local contrast loss does not couple dials. Off-diagonal terms arise only from crosstalk—one dial moving another layer's conditional weight, a violation of dial-locality in the realized apparatus—and are independently measurable by single-dial sweeps. The module's acceptance criteria therefore separate three axes: diagonal deviations within the layer null are apparatus contrast; off-diagonal terms are crosstalk, to be budgeted like any calibration imperfection; only diagonal deviations in excess of the null, with crosstalk budgeted, refute the nested homogeneity premise. The discipline stated above carries over unchanged: the residual-mass law is also the standard quantum-mechanical prediction for ideal sequential measurements (Remark 7.13), so a pass checks consistency, and the empirical force runs in the falsifying direction only.

10 What This Does and Does Not Prove

It proves a conditional bridge. If the effective tracking capacity is Fisher-distinguishability capacity, and if the scalar threshold κ is homogeneous in the physical basis coordinate, then the calibrated binary probability law in θ is Born.

It does not rest on a pending experimental verdict. The bridge is logically downstream of the *operational* reading of IOF—finite-rate basis tracking as reference-frame control. That reading is established classical physics (the data-rate theorem together with phase averaging), consistent with standard quantum mechanics rather than a rival to it. A κ-scaled visibility measurement therefore calibrates an apparatus; it does not adjudicate IOF against quantum mechanics, and it lends no support to the bridge. The open questions are the two bridge assumptions themselves.

It does not derive all of quantum mechanics. The paper does not derive:

- complex Hilbert space;
- full projective quantum geometry;
- tensor-product composition;
- unitary dynamics;
- the Standard Model Lagrangian;
- the ontic admissible-history measure μ_A of IOF;
- the cross-context multi-outcome Born rule: §7 derives the fixed-context form conditionally, but the consistency of weights across incompatible contexts is not derived here.

Those belong to a separate programme—and a companion limit theorem [1] shows that programme cannot be carried out by the same route: the Markov invariance used here generates the Fisher metric and its binary weight but no invariant phase (almost-complex) structure, so relative phase must be imported rather than derived from finite-record geometry. The binary case is the cleanest place where the Born form can be isolated without deriving full Hilbert kinematics.

It removes one degree of freedom from the reconstruction problem. Once finite observer records have Fisher–Rao geometry and scalar-threshold homogeneity holds, the binary form is forced; under nested calibration the same premises force the full weight vector of any single context (§7). The remaining burden is the composition of contexts—phases, projective equivalence, and cross-context consistency—and the exposure of the bridge assumptions to the falsification checks of §9.

It gives the probability law, not the realized record. The squared-cosine form fixes the weights a finite observer must assign to the two records of a calibrated binary context. It does not, and is not meant to, produce *which* record is realized. That step is downstream record formation: the weights are registered as a finite-resolution outcome through the observer's coarse-graining channel, while the single realized record is fixed by the particular admissible history, with IOF's measure μ_A supplying only the ensemble measure [2], not a per-run draw. The law $p(\theta) = \cos^2(\theta/2)$ is therefore the observer's *predictive credence* over records—recovered as a frequency only across the ensemble of histories compatible with the observer's finite information—not an ontic randomizer. The world does not sample; the observer does. The map from this weight to the finite registered record is the coarsening step that belongs to IOF [2] and to the empirical bridge module of the BLQC benchmark [3] (cf. the sufficiency-versus-lossy-coarsening remark above).

11 Objections and Replies

Remark 1 (Is the Born form just a trigonometric identity?). In the binary case, yes: once s is Fisher–Rao arclength, $p(s) = \cos^2(s/2)$ follows by elementary geometry. The substantive claims are upstream: finite observer record invariance selects Fisher–Rao geometry, and scalar-threshold homogeneity selects θ as a Fisher-arclength-affine coordinate.

Remark 2 (Has the Hilbert angle been smuggled in?). No. The angle s is a Fisher–Rao arclength coordinate on a classical probability simplex; θ is the physical basis coordinate of a BLQC experiment. The identification $s = \alpha\theta$ is a scalar-threshold calibration claim, not a Hilbert-space assumption. The two meet only at the level of the calibrated binary record.

Remark 3 (Are the bridge assumptions too strong?). They are strong, but they are not mathematical tautologies. They are empirical claims about how a finite-capacity physical reference channel behaves. If useful tracking capacity is allocated non-uniformly across the basis coordinate, the correct threshold is local, $\kappa(\theta)$, and the derivation fails in that coordinate. The assumptions are exposed as falsifiable calibration claims rather than hidden as definitions.

Remark 4 (Did Wootters already do this?). Wootters showed that statistical distance naturally reproduces the squared-cosine transition-probability form. The present paper asks why a physical observer's finite-rate reference tracking should be calibrated in that geometry at all. The new ingredient is the IOF bridge from useful tracking capacity to Fisher distinguishability, plus scalar-threshold homogeneity of the laboratory coordinate.

Remark 5 (Is the experimental test circular?). The objection assumes the paper makes an experimental discovery claim that the experiment then confirms. It does not. The paper is a conditional foundations derivation: it states two bridge assumptions and derives the binary Born form from them. The associated measurements are methodology, not discrimination. Constant $I(\theta)$ is also a standard-QM prediction, so confirming it confirms quantum mechanics while checking consistency with the bridge; and t_{break} with its dependence on C_{eff} and h_{KS} is a classical control-law (data-rate theorem) prediction consistent with quantum mechanics, which calibrates the apparatus rather than adjudicating between theories. The non-trivial empirical content runs in the falsifying direction: a measured position-dependent $I(\theta)$ in a calibrated record family would reject the derivation (§9).

Remark 6 (Does the choice of hierarchy make the multi-outcome form arbitrary?). The question identifies a real choice, and the paper states it as one. Theorem 7.1 forces only the failure of order-free uniformity; the hierarchical dial-local class is then the motivated arena—it is what a finite-rate channel realizes, binary discriminations a bounded number at a time—and the chain is a canonical representative within it. What is not a choice is the outcome: every member of the class yields the same form (Corollary 7.11). What depends on the choice is the chart and the module's node-mass instance, and those are fixed by the experimenter's own calibration protocol. The choice is protocol; the form is not.

Remark 7 (Is $p_i = q_i^2$ an empty identity?). Pointwise, yes: any distribution is the squared components of some real unit vector, so the identity alone carries no content—and no complex structure is implied by it; the q_i here are real and non-negative, and the companion limit theorem [1] shows this route cannot produce the phase that would make them amplitudes. The content of Theorem 7.10 is the *parametrization* claim: under the stated premises the calibrated physical dials enter as Fisher-arclength coordinates at every node, so the dial dependence of the weights is fixed—and it is exactly this claim, not the pointwise identity, that the Fisher-matrix module measures through the node-mass law.

Remark 8 (Is this already a derivation of quantum mechanics?). No. It is a conditional derivation of the binary probability form in the physical basis coordinate of a BLQC experiment. It does not derive complex Hilbert space, composition, dynamics, or the multi-outcome Born rule. It removes one degree of freedom from the reconstruction problem by showing that, under the stated bridge assumptions, the squared-cosine binary form is forced.

12 Related Work

The relation between statistical distance and quantum transition probabilities is not new. Wootters [8] showed that quantum distinguishability is naturally expressed in terms of statistical distance. This paper uses the logic in the opposite direction. It starts from finite observer record constraints, invokes the Cencov uniqueness of Fisher–Rao geometry under sufficient Markov morphisms, and obtains the squared-coordinate binary form from the resulting record geometry. The IOF bridge then identifies the laboratory basis coordinate as the Fisher-arclength-affine coordinate via the scalar-threshold reading. The same attribution governs the multi-outcome extension: the statistical-distance identity on the n-outcome simplex is Wootters's, and the conditional (stick-breaking) coordinates with their Fisher–Rao orthogonality are classical [6]; §7 contributes the obstruction theorem, the hierarchical calibration reading, and the node-mass module.

The novelty claim should remain narrow. The paper does not claim to be the first observation that Fisher geometry and the Born rule are related. The proposed contribution is the directional chain

$$\begin{gathered}\text{IOF finite-rate basis tracking}\\ \longrightarrow\\ \text{Fisher capacity bridge}\\ \longrightarrow\\ \text{Cencov: Fisher–Rao on record manifold}\\ \longrightarrow\\ \text{square-root coordinates: } p(s)=\cos^2(s/2)\\ \longrightarrow\\ \text{scalar-threshold homogeneity: } s=\alpha\theta\\ \longrightarrow\\ p(\theta)=\cos^2(\theta/2) \text{ in the calibrated lab basis coordinate.}\end{gathered}$$

This should also be distinguished from broader quantum reconstructions, including those that derive Hilbert-space structure from operational axioms [9, 10], purification principles, or agent-centred probability [11]. Those programmes aim at the full kinematics of quantum theory. The present paper isolates only the binary probability form available once the observer's record geometry has been fixed and the lab basis coordinate has been pinned to it via IOF.

A different route to the Born weights is Zurek's derivation from *envariance* [12], which obtains $p_k = |\psi_k|^2$ from a symmetry—invariance under entanglement-assisted swaps—already operating inside complex Hilbert space. That derivation and the present one reach the same weight from opposite ends. Envariance descends from the full entangled Hilbert-space structure; the route here ascends from classical finite-record (Fisher–Rao) geometry and stops before phase, composition, and entanglement—provably so, by the companion limit theorem [1]. The two are therefore complementary rather than competing: envariance presupposes exactly the structure this paper declines to assume, while this paper isolates the binary weight without it.

13 Conclusion

The binary Born form in the laboratory basis coordinate θ can be read as the expression of Fisher–Rao record geometry in square-root coordinates for a calibrated finite record, with the lab coordinate pinned to the Fisher arclength by scalar-threshold homogeneity. On this reading, the squared cosine is not a separate quantum probability postulate at the level of the binary case. It is the coordinate form forced by invariant finite-record distinguishability together with the Fisher capacity bridge.

The conditional weight is carried by two named assumptions: the Fisher capacity bridge (Assumption 2.1) and scalar-threshold homogeneity (Assumption 5.1). Both are empirically meaningful, and the Fisher-homogeneity module of the BLQC benchmark [3] engages them in the only direction that carries force: a measured position-dependent $I(\theta)$ in a calibrated record family would reject the derivation here. Confirmation is weaker by nature—constant $I(\theta)$ is also a standard-QM prediction, so the module checks consistency rather than singling the bridge out, and κ-scaled visibility loss, being classical control physics, lends the bridge no support at all.

What is no longer a separate problem is the identification of the Fisher arclength coordinate with the laboratory basis angle. The claim is correspondingly interpretive, not rival: IOF does not replace the Born rule; it interprets the binary Born rule as the record-level form taken by finite directional sampling of an underlying quantum state.

The extension of §7 carries the same reading to the full weight vector of a single calibrated context: the curvature of record geometry rules out any uniform multi-dial calibration, hierarchical binary discrimination is the motivated response, and under the conditional form of the same two assumptions the fixed-context Born form—squared components of a real unit vector, dial-affine in Fisher arclength at every node—follows for every n and every calibration topology in the class. What this route provably does not reach—the composition of incompatible contexts and the phase that governs it [1]—marks the boundary between record geometry and quantum structure.

References

[1] Aernoud Dekker. Why fisher geometry gives binary born but not quantum phase: A limitative theorem on markov-invariant record geometry, 2026. OSF preprint.

[2] Aernoud Dekker. The ignorant observer: Avidya as a physical limit, 2026. OSF Preprints.

[3] Aernoud Dekker. Bandwidth-limited quantum control: A finite-rate phase-reference control law and its operational benchmark, 2026. OSF preprint.

[4] N. N. Cencov. *Statistical Decision Rules and Optimal Inference*, volume 53 of *Translations of Mathematical Monographs*. American Mathematical Society, 1982.

[5] Shun-ichi Amari and Hiroshi Nagaoka. *Methods of Information Geometry*, volume 191 of *Translations of Mathematical Monographs*. American Mathematical Society, 2000.

[6] Robert J. Connor and James E. Mosimann. Concepts of independence for proportions with a generalization of the Dirichlet distribution. *Journal of the American Statistical Association*, 64(325):194–206, 1969.

[7] Stephen D. Bartlett, Terry Rudolph, and Robert W. Spekkens. Reference frames, superselection rules, and quantum information. *Reviews of Modern Physics*, 79:555–609, 2007. doi: 10.1103/RevModPhys.79.555.

[8] William K. Wootters. Statistical distance and hilbert space. *Physical Review D*, 23(2): 357–362, 1981.

[9] Lucien Hardy. Quantum theory from five reasonable axioms, 2001. arXiv:quant-ph/0101012.

[10] Giulio Chiribella, Giacomo Mauro D'Ariano, and Paolo Perinotti. Informational derivation of quantum theory. *Physical Review A*, 84(1):012311, 2011.

[11] Christopher A. Fuchs, N. David Mermin, and Rüdiger Schack. An introduction to QBism with an application to the locality of quantum mechanics. *American Journal of Physics*, 82(8):749–754, 2014.

[12] Wojciech H. Zurek. Probabilities from entanglement, Born's rule $p_k = |\psi_k|^2$ from envariance. *Physical Review A*, 71(5):052105, 2005.