The Born Rule from Finite Observation

A Conditional Derivation of the Binary Born Form

Aernoud Dekker

June 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 is conditional on both 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.

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

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

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*

$$g^{(m)}_{Kp}(Ku, Kv) = g^{(n)}_p(u, v)$$

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 $\kappa = h_{\mathrm{KS}} - C_{\mathrm{eff}} \ln 2$ 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 §8.

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 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

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

—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,

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

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

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

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 [6]—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.

8 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 $\kappa_{\mathrm{cal}} = h_{\mathrm{KS}} - \eta\, C_{\mathrm{eff}} \ln 2$ 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 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 multi-outcome Born rule for arbitrary projective measurements.

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. The remaining burden is to derive the multi-outcome rule under additional structure (phases, projective equivalence, composition, context constraints), and to expose the bridge assumptions to the falsification check of §8.

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).

10 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 (§8).

Remark 6 (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.

11 Related Work

The relation between statistical distance and quantum transition probabilities is not new. Wootters [7] 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 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

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

This should also be distinguished from broader quantum reconstructions, including those that derive Hilbert-space structure from operational axioms [8, 9], purification principles, or agent-centred probability [10]. 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* [11], 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.

12 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.

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] 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.

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

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

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

[10] 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.

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