Why Fisher Geometry Gives Binary Born but Not Quantum Phase

A Limitative Theorem on Markov-Invariant Record Geometry

Aernoud Dekker

June 2026

A companion to *The Born Rule from Finite Observation*

Abstract

The companion paper *The Born Rule from Finite Observation* [1] derives the binary Born weight $p(\theta) = \cos^2(\theta/2)$ from the Fisher–Rao geometry of finite observer records, under Markov invariance plus a bandwidth-limited control (BLQC) bridge. This paper establishes the matching *limit*. Using the Chentsov–Campbell classification and its arbitrary-degree extension by Ay, Jost, Lê, and Schwachhöfer—that the only covariant 2-tensor field on finite probability simplices invariant under sufficient Markov morphisms is, up to scale, the symmetric Fisher–Rao metric—we show that classical finite-record geometry admits *no invariant almost-complex structure at all*: the only Markov-invariant $(1,1)$-tensor is a scalar multiple of the identity, and no real scalar multiple of the identity squares to $-\mathrm{Id}$. Consequently the relative phase that distinguishes quantum amplitudes from probability vectors—and with it the symplectic (Kähler) structure that governs interference between non-commuting contexts—is not generated as an invariant tensorial structure of finite-record geometry. It must be imported, or derived from additional non-classical structure such as composition or local tomography. The result is a *generation* theorem, not a *reproduction* theorem: it does not claim that classical models cannot reproduce quantum statistics, only that phase-sensitive context structure is not an invariant of classical finite-record geometry. Binary Born is exactly where Fisher geometry reaches; quantum phase is exactly where it stops.

1 Purpose

The companion paper [1] shows that, for a single calibrated two-outcome measurement, the Born weight $p(\theta) = \cos^2(\theta/2)$ is not a separate probability postulate. It is the coordinate form forced by the invariant distinguishability geometry of finite observer records—the Fisher–Rao metric selected by Chentsov's theorem—once a BLQC bridge pins the laboratory basis angle to the Fisher arclength. That is a positive result with a precise reach: one calibrated binary context.

It is natural to ask whether the same machinery extends to the full, multi-outcome, multi-context Born rule $P(\Pi) = \langle\psi|\Pi|\psi\rangle$. This paper answers in the negative, and makes the negative answer sharp. The obstruction is not a missing calculation or a consistency theorem yet to be proved. It is structural: the multi-outcome Born rule lives on complex projective space, whose defining extra ingredient over a classical probability simplex is *relative phase*—the source of interference between non-commuting measurement contexts. We show that relative phase, in its differential-geometric form as an almost-complex (equivalently, symplectic) structure, is *not generated* by any Markov-invariant construction on finite record distributions. The very invariance that selects Fisher–Rao and powers the binary derivation is what forbids the phase structure required for the rest of quantum kinematics.

The contribution is therefore to convert an obstruction into a theorem, and to delimit exactly where finite-record information geometry has explanatory power over quantum probability. The boundary is clean:

$$\boxed{\text{Binary Born} = \text{the reach of Fisher geometry}; \qquad \text{quantum phase} = \text{where it stops.}}$$

This does not shrink the binary result; it locates it. The binary weight is the maximal Born content extractable from classical record invariance, and the limit theorem certifies that nothing past it comes for free.

2 Setup: Records, Simplices, and the Invariance Class

A finite observer's accessible record is a probability vector on a finite alphabet, i.e. a point of the open simplex

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

Its tangent space at p is $T_p\Delta_n^\circ = \{u \in \mathbb{R}^n : \sum_i u_i = 0\}$. This is the same operational object used in [1]: not yet a quantum state, but a finite distribution over distinguishable record outcomes.

The admissible transformations of records are the *sufficient Markov morphisms*: stochastic maps that change the representation of a record without losing information about the tracked family. These are exactly the maps under which the companion paper requires distinguishability to be preserved [1, Assumption 2]. Concretely they are generated by

1. **relabelings** — bijective permutations of outcomes (deterministic, invertible, hence sufficient), including *odd* permutations;
2. **congruent embeddings (sufficient refinements)** — splitting an outcome i into several with fixed, θ-independent branching ratios, $p_i \mapsto (\lambda_1 p_i, \ldots, \lambda_k p_i)$, and their left inverses (sufficient coarse-grainings).

Definition 1 (Markov-invariant tensor field)**.** *A family $\{T^{(n)}\}$ of covariant tensor fields, one on each Δ_n°, is* Markov-invariant *(a* congruent family*) if every sufficient Markov morphism $f : \Delta_n^\circ \to \Delta_m^\circ$ pulls the higher tensor back to the lower one, $f^*T^{(m)} = T^{(n)}$, on the model directions. The same definition applies to $(1,1)$-tensor fields (endomorphism fields of the tangent bundle), with pullback replaced by the natural transport induced by f.*

Two points about this invariance class matter for what follows, and both are honest hypotheses rather than conveniences.

Remark 1 (The full symmetric group is required)**.** *The class includes odd permutations and dimension-changing embeddings, not merely an orientation-preserving subgroup. This is essential: under a weaker, orientation-preserving group an even-dimensional simplex* would *admit an invariant volume form, and the limitative theorem below would fail. Requiring invariance under the full set of sufficient morphisms is precisely the hypothesis already in force in [1], so the limit theorem and the positive result rest on the* same *premise; see Remark 3 and §7.*

Remark 2 (Same invariance, both directions)**.** *Chentsov's theorem uses this invariance to* select *a unique metric; we use the same invariance to* forbid *a phase structure. One classification theorem, read at its symmetric and antisymmetric parts, yields both the reach (§3) and the limit (§5).*

3 The Reach: Fisher–Rao and the Binary Born Weight

We recall the positive result in geometric form. By Chentsov's theorem [2], extended to the simplex by Campbell [3] (modern treatment [4]), the unique-up-to-scale Markov-invariant Riemannian metric on Δ_n° is Fisher–Rao,

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

In square-root coordinates $q_i = \sqrt{p_i}$ this is the round metric: $g = 4\sum_i dq_i^2$, so (Δ_n°, g) is isometric to the open positive orthant of the unit sphere $S^{n-1} \subset \mathbb{R}^n$, a *real*, totally geodesic submanifold. For a binary record $p = (p, 1-p)$, the Fisher arclength $s = 2\arccos\sqrt{p}$ gives the identity

$$p(s) = \cos^2(s/2), \qquad 1 - p(s) = \sin^2(s/2),$$

and the BLQC scalar-threshold bridge of [1] fixes $s = \theta$ on the calibrated interval, yielding $p(\theta) = \cos^2(\theta/2)$. This is the reach: a one-real-parameter law along a geodesic of the real Fisher sphere.

The square root here is not a quantum amplitude; it is the coordinate induced by the invariant statistical metric, and the geometry it produces is real-orthogonal. The question of the next section is whether the *complex* amplitude—the part of a quantum state that a probability vector does not record—can be recovered from the same invariant geometry. It cannot.

4 The Classification Input

The limit theorem rests on one imported result, which we state in the form we need. Let τ_k denote the *canonical k-tensor* on statistical models,

$$\tau_k(V_1, \dots, V_k) = \mathbb{E}_p[\partial_{V_1} \log p \cdots \partial_{V_k} \log p]\,,$$

so that τ_2 is the Fisher metric and τ_3 is the Amari–Chentsov tensor. Each τ_k is *symmetric* in its arguments, because the integrand is a product of scalars and multiplication commutes.

Theorem 1 (Chentsov–Campbell–AJLS classification)**.** *On the finite simplices Δ_n°, every Markov-invariant covariant tensor field is algebraically generated by the canonical tensors $\{\tau_k\}$ [5, 6]. In particular, the space of Markov-invariant covariant 2-tensor fields is one-dimensional, spanned by the Fisher metric $\tau_2 = g$.* (Classical cases: [2, 3]; arbitrary degree without a symmetry hypothesis: [5, 6]; survey: [7].)

The invariance hypothesis of Theorem 1 is exactly Assumption 2 of the companion paper [1]—that admissible finite-resolution changes are sufficient Markov maps preserving distinguishability of the tracked family. The positive binary result and the negative phase result therefore rest on one and the same premise, read at the symmetric and antisymmetric parts of the invariant tensor algebra (cf. Remark 2).

The crucial feature for us is that Theorem 1 does *not* presuppose symmetry: it classifies all covariant 2-tensors, symmetric or not, and finds the invariant ones exhausted by the symmetric Fisher metric. Hence:

Corollary 1 (No invariant antisymmetric 2-tensor)**.** *There is no nonzero Markov-invariant antisymmetric covariant 2-tensor (no invariant 2-form) on the finite simplices. Equivalently, the antisymmetric part of any Markov-invariant 2-tensor vanishes.*

Proof. By Theorem 1 an invariant 2-tensor is a scalar multiple of g, which is symmetric; its antisymmetric part is therefore zero. □

Remark 3 (Hand check at $n = 2, 3$)**.** *For $n = 2$, $\dim \Delta_2^\circ = 1$ and there is no nonzero antisymmetric 2-tensor at all—which is exactly why the binary case is unobstructed. For $n = 3$, $\dim \Delta_3^\circ = 2$, and an antisymmetric 2-tensor is a multiple $f(p)\, dA$ of an area form. A transposition of two outcomes is a sufficient Markov morphism that reverses orientation, $dA \mapsto -dA$; invariance under the full symmetric group then forces $f \equiv 0$. The Fisher metric, being symmetric, survives the same transposition. This is Remark 1 made concrete: the parity-odd morphisms are what kill the candidate symplectic form.*

We also need the closure of invariance under index raising, which is standard and is noted explicitly in the Markov-invariance literature [7].

Lemma 1 (Invariant $(1, 1)$-tensors)**.** *Since the Fisher metric g is itself Markov-invariant and nondegenerate, lowering an index with g is a natural isomorphism between invariant $(1, 1)$-tensor fields and invariant $(0, 2)$-tensor fields. Hence the space of Markov-invariant $(1, 1)$-tensor fields is one-dimensional, spanned by the identity endomorphism* Id.

Proof. If J is an invariant $(1, 1)$-tensor, then $\tilde{J}(u, v) := g(Ju, v)$ is an invariant $(0, 2)$-tensor: because g is itself Markov-invariant and nondegenerate, the musical isomorphism (index lowering by g) commutes with the transport induced by every sufficient morphism, so the invariance of J transfers to $\tilde{J}$. By Theorem 1, $\tilde{J} = c\, g$ for some scalar c; nondegeneracy of g gives $J = c\,\mathrm{Id}$. Conversely Id is invariant. The map $J \mapsto \tilde{J}$ is bijective because g is nondegenerate. □

5 The Limit: No Invariant Almost-Complex Structure

An *almost-complex structure* is a $(1, 1)$-tensor field J with $J^2 = -\mathrm{Id}$. It is the differential-geometric carrier of a relative phase: it rotates a tangent direction into its phase-conjugate partner, and through $\omega(u, v) = g(Ju, v)$ it furnishes the symplectic form whose nonvanishing is interference. A quantum pure-state space $\mathbb{CP}^{n-1}$ is Kähler: it carries a compatible triple $(g_{\mathrm{FS}}, J, \omega_{\mathrm{FS}})$, and ω_{FS} is the structure that makes amplitudes interfere [9, 10].

Theorem 2 (No invariant almost-complex structure on finite records)**.** *The finite record simplices Δ_n° admit no nonzero Markov-invariant almost-complex structure. Consequently they admit no Markov-invariant Kähler or symplectic structure, and in particular none compatible with the Fisher–Rao metric.*

Proof. The argument is six lines.

1. On Δ_n°, sufficient Markov invariance leaves the space of covariant 2-tensors spanned by Fisher–Rao (Theorem 1).

2. Hence every invariant $(0, 2)$-tensor is symmetric (Corollary 1).

3. Since g is invariant and nondegenerate, an invariant $(1, 1)$-tensor J yields the invariant $(0, 2)$-tensor $g(J\cdot, \cdot)$ (Lemma 1).

4. By steps 1–2 that tensor equals $c\, g$, so $g(JX, Y) = c\, g(X, Y)$ for all X, Y, i.e. $J = c\,\mathrm{Id}$.

5. $J^2 = -\mathrm{Id}$ then requires $c^2 = -1$, impossible for a real scalar acting on real tangent spaces.

6. Therefore no invariant J with $J^2 = -\mathrm{Id}$ exists.

A Markov-invariant symplectic or Kähler structure compatible with g would supply such an invariant J via $J = g^{-1}\omega$, so no invariant such structure exists either. An arbitrary symplectic form chosen by hand on an even-dimensional simplex is not excluded—and is not the point; only the absence of a *natural, Markov-invariant* one is claimed. □

The strength of the statement is worth marking. We do *not* merely show that no almost-complex structure compatible with the Fisher metric exists; we show that no invariant almost-complex structure exists *at all*, compatible or not, because the only invariant $(1, 1)$-tensor is $c\,\mathrm{Id}$. The Fisher-compatible (Kähler) case is the special instance $c\,\mathrm{Id}$ with c chosen to be metric-orthogonal, and it is excluded along with the rest. There is no second invariant endomorphism to play the role of multiplication by i.

Remark 4 (Why the binary weight is still reached)**.** *At $n = 2$ the obstruction in Theorem 2 is automatic for a trivial reason—an odd-dimensional manifold carries no almost-complex structure—while the substantive parity obstruction of Remark 3 first bites at* $n = 3$. *Either way, the qubit's amplitude phase is already outside the reach of finite-record geometry; see Corollary 2. The binary Born* weight *is reached precisely because it is the projection onto the one direction (the modulus) that the simplex does record.*

6 Corollary: Relative Phase Is Not Generated

The geometric theorem has a direct physical reading through the modulus map. Fix a basis and let

$$\pi : \ \mathbb{CP}^{n-1} \longrightarrow \overline{\Delta_n}, \qquad [\psi] \longmapsto \big(|\langle i|\psi\rangle|^2\big)_{i=1}^{n}$$

be the modulus-squared projection of a pure state to its outcome distribution—the probability-forgetting map that, in quantum theory, returns the Born probabilities. Counting real dimensions,

$$\dim_{\mathbb{R}} \mathbb{CP}^{n-1} = 2(n-1) = \underbrace{(n-1)}_{\text{moduli (the simplex)}} + \underbrace{(n-1)}_{\text{relative phases (the fibre)}}.$$

The classical record is the *base* of this fibration; the relative phases are the *fibre*—over the interior Δ_n° a torus T^{n-1} of $U(1)$ factors, degenerating on the boundary where zero-amplitude components carry no phase—quotiented away the moment one writes down a probability vector. Phase information is therefore absent from the definition of a finite record, before any geometry is imposed.

Corollary 2 (Phase is import-only)**.** *The relative-phase data distinguishing a quantum amplitude from its outcome distribution—equivalently, the Kähler form ω_{FS} on $\mathbb{CP}^{n-1}$ and the interference between non-commuting contexts it governs—does not descend to, and is not generated as a Markov-invariant tensorial structure on, the finite record simplex. Any reconstruction that recovers the multi-outcome, multi-context Born rule must* import or independently derive *that structure; finite-record information geometry supplies it for no value of n, including the qubit $n = 2$.*

Proof. The fibre of π carries the phase data and is collapsed by passage to a record, so phase is not a function on the simplex. Theorem 2 adds the differential statement: the simplex carries no invariant almost-complex structure that could re-encode a phase tangentially. Both the value data and the geometric carrier are therefore absent from invariant record geometry. □

It is worth naming the causal order. The absence of an invariant almost-complex structure (Theorem 2) is a *symptom*; the root cause is prior and cheaper—passing to a probability vector has already collapsed the phase fibre, so no phase remains for any tangential structure to re-encode. Theorem 2 is the differential shadow of a collapse already built into the definition of a record.

Thus the reach and the limit are two corollaries of one classification. The symmetric part of the invariant 2-tensor algebra gives the Fisher metric and the binary Born weight (§3); the vanishing of the antisymmetric part gives the absence of phase (§5). Binary Born is the projection onto the modulus, which the record keeps; quantum phase is the fibre, which the record forgets.

7 What This Does and Does Not Claim

The theorem is easy to over-read in two opposite directions. We fence it explicitly.

It is a generation theorem, not a reproduction theorem. We do *not* claim that classical models cannot reproduce quantum statistics. They can: any generalized-probabilistic state space embeds in a large enough simplex, ontological models exist, and the statistics of a single fixed context are trivially classical. The claim is the strictly weaker, and correct, one: phase-sensitive quantum context structure is not generated as an *invariant tensorial structure of classical finite-record geometry*. What can be encoded by hand or by extra postulates is a separate question; what the invariant geometry produces on its own is the Fisher metric and nothing more.

The full symmetric group is a load-bearing hypothesis. The obstruction needs invariance under odd permutations and congruent embeddings (Remarks 1, 3). Under a merely orientation-preserving group, even-dimensional simplices admit an invariant volume form and the theorem fails. This hypothesis is not an addition: it is exactly the sufficiency/invariance premise already used to *select* Fisher–Rao in the companion derivation, so the positive and negative results share one premise rather than trading premises.

It is silent on complex versus real versus quaternionic. The theorem generates nothing beyond the metric, so it does not on its own say "you need *complex* Hilbert space"—only that you must import *some* phase/context structure. It therefore does not distinguish complex from real or quaternionic quantum theory; which additional structure is selected depends on further composition principles such as local tomography [13, 14, 15]. The present result locates the import; it does not perform it.

It concerns invariant tensor fields. The result is about geometric structure invariant under sufficient morphisms. It does not forbid introducing phase through additional, non-tensorial primitives—extra postulated structure, a chosen connection, or a composition rule. That is precisely the "must import" conclusion, not a gap in it. The α-connections of information geometry, for instance, are invariant but are affine connections, not $(1, 1)$-tensors, and do not furnish a J; they do not rescue phase.

8 Consequence for the Reconstruction Programme

Three consequences follow for any attempt to push the finite-record programme past the binary case.

The binary result is non-redundant exactly where Gleason is silent. Gleason's theorem [11] fixes every Born probability, including all binary marginals, once the projector lattice of a Hilbert space of dimension ≥ 3 and noncontextuality are granted; there the finite-record binary derivation adds nothing. Standard Gleason is silent in dimension two, and the POVM strengthening [12] restores it only by assuming the full complex effect algebra. The finite-record binary weight is thus load-bearing precisely in the $d = 2$ gap that the lattice-based theorems cannot reach unaided—and Theorem 2 explains why that is the natural ceiling, not a temporary one.

The missing object is the space, not the formula. The next theorem worth proving is not "binary coarse-grainings imply Born." Any such statement that mentions projective alternatives, subspaces, or noncontextual measures has already imported the Hilbert lattice, after which Gleason or Busch does the work and the binary input is supportive at best. The genuine open problem is the prior one: *why do admissible observer records organize as the contexts of a* *complex projective space at all?*—that is, why local tomography, reversible/continuous context transformations, and tensor composition. Theorem 2 certifies that none of these can be read off the invariant geometry of the records; they are independent structure.

The only finite-record-native route to that structure is dynamical. The operational reconstructions [13, 14, 15] already derive complex Hilbert space from composition-type axioms; restating them in record vocabulary adds nothing. The one ingredient the Ignorant Observer Framework possesses that those programmes do not is a *finite-capacity* observer with a breakdown threshold $\kappa = h_{\mathrm{KS}} - C_{\mathrm{eff}} \ln 2$. A strong continuation is worth attempting only if finite capacity supplies an *independent physical reason* for local tomography or for reversible composition—if it *forces* the composition axiom rather than merely tolerating it. Absent that, Theorem 2 marks the honest endpoint: publish the binary reach, certify the phase limit, and treat the complex-context structure as imported.

9 Objections and Replies

Remark 1 (Is this just "classical is not quantum"?)**.** *No. That slogan is about reproducibility of statistics and is false as stated (simplex embeddings reproduce any single-context statistics). The result here is a precise* generation *statement located at a named obstruction—the vanishing of the invariant antisymmetric 2-tensor—and it is used to delimit a specific positive claim, the binary reach, rather than to gesture at a difference. The novelty is the location of the obstruction and its identity with the Chentsov classification, not the existence of a difference.*

Remark 2 (Could the area form on Δ_3 be the missing symplectic form?)**.** *It is invariant only under orientation-preserving relabelings. A single transposition of two outcomes—an admissible sufficient morphism—reverses it (Remark 3). Under the full invariance class it is not invariant, and there is no other candidate. This is exactly why the hypothesis must name the full symmetric group.*

Remark 3 (Quantum state space is Kähler, so phase *is* compatible with a metric—doesn't that contradict the theorem?)**.** *No; it sharpens it. On $\mathbb{CP}^{n-1}$ the triple $(g_{\mathrm{FS}}, J, \omega_{\mathrm{FS}})$ is compatible, and g_{FS} restricted to the real (phase-zero) slice is the Fisher round metric. The theorem says the* other *two members of the triple, J and ω_{FS}, are not invariants of the record geometry. The Kähler case is not an escape; it is the special compatible case that Theorem 2 excludes along with every other J.*

Remark 4 (Does the g-isomorphism step secretly assume what it proves?)**.** *No. It assumes only that g is invariant and nondegenerate—both standard—so that raising and lowering indices commute with morphism transport [7]. The classification of invariant $(0, 2)$-tensors is then imported once (Theorem 1) and transported to $(1, 1)$-tensors. No complex or phase structure is used anywhere in the argument; that is the point.*

Remark 5 (Is the result merely an artefact of demanding strict invariance instead of monotonicity?)**.** *Strict invariance is the correct notion for* sufficient *morphisms (no information loss), which is the regime in which Fisher–Rao is selected in the first place [2, 3, 5]. Lossy, merely monotone maps form a separate (contractive) story and do not generate new invariant structure; if anything they destroy structure. The premise here is identical to the one that makes the binary derivation go through.*

10 Relation to the Binary Paper

The two papers are one argument read at two parities. The companion [1] uses Chentsov invariance to select the symmetric Fisher metric and, with the BLQC bridge, extracts the binary

Born weight from it. The present paper uses the *same* invariance, applied to the antisymmetric and $(1,1)$ parts, to certify that no phase structure accompanies the metric. Neither result weakens the other. Together they state a boundary with no interior gap: everything the invariant geometry yields is the Fisher metric and its binary Born weight; everything quantum past that—phase, interference, non-commuting contexts, composition—is provably outside it and must be hosted or derived elsewhere.

11 Conclusion

Finite-record information geometry reaches the binary Born weight and stops there. The reach is the symmetric content of the unique Markov-invariant 2-tensor; the limit is the vanishing of its antisymmetric content, which forbids any invariant almost-complex structure and with it the relative phase that drives quantum interference. The boundary is exact and falls out of a single classification theorem. For the broader programme this is clarifying rather than discouraging: the binary derivation is not a fragment of an unfinished full-Born proof but the complete account of where classical record invariance has purchase, and the limit theorem certifies that the remaining quantum structure—phase, composition, the complex projective context space—is genuinely additional, to be imported under a separate principle (local tomography) or, if the Ignorant Observer Framework is to contribute something new, to be forced by finite capacity itself. Until such a forcing is exhibited, binary Born is exactly where Fisher geometry gives the Born rule, and quantum phase is exactly where it does not.

References

[1] A. Dekker, *The Born Rule from Finite Observation: A Conditional Derivation of the Binary Born Form*, OSF Preprints (2026). DOI: 10.17605/OSF.IO/U5RDE.

[2] N. N. Cencov (Chentsov), *Statistical Decision Rules and Optimal Inference*, Translations of Mathematical Monographs, Vol. 53, American Mathematical Society, 1982.

[3] L. L. Campbell, An extended Čencov characterization of the information metric, *Proceedings of the American Mathematical Society* 98 (1986), 135–141.

[4] S. Amari and H. Nagaoka, *Methods of Information Geometry*, Translations of Mathematical Monographs, Vol. 191, American Mathematical Society, 2000.

[5] N. Ay, J. Jost, H. V. Lê, and L. Schwachhöfer, Information geometry and sufficient statistics, *Probability Theory and Related Fields* 162 (2015), 327–364. arXiv:1207.6736.

[6] N. Ay, J. Jost, H. V. Lê, and L. Schwachhöfer, Congruent families and invariant tensors, in *Information Geometry and Its Applications* (IGAIA IV), N. Ay, P. Gibilisco, and F. Matúš (eds.), Springer Proceedings in Mathematics & Statistics, Vol. 252, Springer, 2018, pp. 157–187. DOI: 10.1007/978-3-319-97798-0_6. arXiv:1705.11014.

[7] A. Fujiwara, Hommage to Chentsov's theorem, *Information Geometry* 7 (Suppl. 1) (2024), 79–98. DOI: 10.1007/s41884-022-00077-7.

[8] W. K. Wootters, Statistical distance and Hilbert space, *Physical Review D* 23 (1981), 357–362.

[9] T. W. B. Kibble, Geometrization of quantum mechanics, *Communications in Mathematical Physics* 65 (1979), 189–201.

[10] A. Ashtekar and T. A. Schilling, Geometrical formulation of quantum mechanics, in *On Einstein's Path*, Springer, 1999, 23–65. arXiv:gr-qc/9706069.

[11] A. M. Gleason, Measures on the closed subspaces of a Hilbert space, *Journal of Mathematics and Mechanics* 6 (1957), 885–893.

[12] P. Busch, Quantum states and generalized observables: a simple proof of Gleason's theorem, *Physical Review Letters* 91 (2003), 120403.

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

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

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