

The Squaring Residue

Why Reversible Probability-Preserving Dynamics Selects ℓ^2 — and the Single Substrate Question on Which It Turns

Keith Taylor VERSF Theoretical Physics Programme versf-eos.com

For the General Reader

Quantum mechanics turns an amplitude into a probability by *squaring* its size: $P = |\psi|^2$. A century of work has not explained why the exponent is 2 and not something else. The preceding papers in this programme reduced almost everything about the quantum "stage" to a more primitive counting principle, and in doing so isolated the last genuinely unexplained ingredient — the squaring itself — and showed it was an old, named open problem (PAMV Open Problem 10). This paper attacks that single residue.

The natural temptation is to try to *prove* that the squared length is forced. We resist it, because every proof of that kind turns out to assume its own conclusion: it defines the allowed motions of the theory as "the ones that preserve a squared length," and then announces, unsurprisingly, that they preserve a squared length. That is not a derivation; it is the answer written into the question.

Instead we do something more honest and, it turns out, more decisive. We refuse to start from the allowed motions and start from the *conservation law* — the requirement that whatever bookkeeping quantity the substrate conserves is additive over outcomes you can tell apart. From that alone, plus the phase structure the earlier papers already established, the conserved quantity is pinned down to a simple separable form: a sum, over outcomes, of some fixed function of each amplitude's size. Crucially, *every* exponent p survives this stage — the squaring is not yet forced, and that is correct: it should not be forcible from conservation and phase alone.

What forces the squaring is one extra thing, and we can name it exactly. There is a clean, old mathematical fact: for any exponent other than 2, the continuous symmetries that preserve a "sum of p -th powers" are *only* the ones that rotate each outcome's phase independently — they never mix two distinct outcomes into each other. Only at $p = 2$ does the symmetry enlarge to include continuous *mixing* between distinct outcomes — the beam-splitter-like motions at the heart of quantum interference between alternatives. So the entire question "why is the exponent 2?" collapses to a single, concrete, physical question about the substrate:

Does the substrate's reversible dynamics include continuous motions that coherently redistribute amplitude *between distinct outcomes* — not merely rotate each outcome's phase in place?

If yes, the exponent is 2, full stop, and the squaring is derived. If the substrate's reversible motions are only phase rotations of individual outcomes (plus discrete relabellings), then every

exponent is consistent and the squaring is a genuinely independent fact about reality. Either answer settles a century-old question, and either is worth having. This paper does not claim to know which holds; it proves that this is the *only* thing left to know, and that it is a sharp, answerable, substrate-level question rather than a vague philosophical one.

Abstract

The ODG-reduction paper reduced Operational Distinguishability Geometry to finite distinguishability packing together with a path-sum substrate (H0), and localized the one component that does not descend constructively — the *inner product* of the carrier — to a single residue: that the conserved, reversible-transport-invariant normalization on the outcome-amplitude vector \mathbb{C}^d is the squared (ℓ^2) norm. That residue was shown identical to PAMV Open Problem 10 and given three interderivable faces (U(d)-invariance of the 2-norm \equiv positive-semidefiniteness of the correlation kernel \equiv bilinearity of selection). The present paper attacks the residue directly.

We decline the affirmative framing — "prove the reversible group is U(d)" — on the ground established in the prior paper: every route to it presupposes "reversible = isometry of a complex inner product," which is the conclusion. We adopt instead the **conservation-first** framing recommended by the prior paper's own §11.1, and we lead with bit conservation (BCB) on the *normalization functional*, not with the symmetry group.

The analysis yields a sharp dichotomy. Three results are **proven** (modulo inherited inputs):

1. **(Separability from BCB.)** BCB additivity over disjoint outcomes, together with the inherited diagonal U(1) per-path phase conservation, forces the conserved normalization to be a separable, phase-blind functional $N(\psi) = \sum_i h(|c_i|)$ in any outcome basis — for *some* function h , with no exponent yet selected.
2. **(The Diagonal-Torus Lemma.)** For every $p \neq 2$ ($1 \leq p < \infty$), the connected component of the ℓ^p -isometry group on \mathbb{C}^d ($d \geq 2$) is exactly the diagonal phase torus $U(1)^d$; only at $p = 2$ does *the norm's invariance group* enlarge to the full unitary group U(d). (This is a statement about the symmetry group G_N of the functional, not about how much of it the substrate dynamics R fills — a distinction §5 leans on: a single off-diagonal generator forces $N = \ell^2$, but $R^0 = U(d)$ requires generators connecting all coordinate pairs.) Equivalently: a continuous one-parameter group of N -preserving transformations that acts *off-diagonally* — coherently mixing the amplitudes of two distinct outcomes — exists if and only if N is (up to normalization) the ℓ^2 norm.
3. **(The Reversible-Mixing Dichotomy.)** Combining 1 and 2: the squaring residue resolves affirmatively ($p = 2$, the inner product is recovered, PAMV OP10 closes positively) **iff** substrate-admissible reversible dynamics contains a continuous generator that coherently redistributes amplitude between distinct admissible outcomes. If reversible dynamics is exhausted by the diagonal phase torus (per-path U(1)) together with discrete outcome relabellings, then every ℓ^p is substrate-admissible, the inner product is independent structure, and PAMV OP10 settles negatively.

The residue is thereby relocated from a condition on an abstract symmetry group to a single concrete physical question — **does the substrate support continuous reversible amplitude-redistribution between distinct outcomes?** — which is interpretable, falsifiable, and decidable at substrate level. What the paper does **not** do is decide it: whether TPB/BCB plus temporal extensibility *force* such an off-diagonal generator is left open, with a structured argument for why the inherited per-path $U(1)$ holonomy is *insufficient* to supply it (it lives in the diagonal torus and is consistent with all p) and a sketch of what an affirmative substrate argument would have to establish (continuous coherent inter-outcome transport, i.e. a substrate beam-splitter). The contribution is the dichotomy and its sub-Hilbertian proof: "why ℓ^2 " is exactly "is there continuous off-diagonal reversible mixing," and not one assumption more.

Contents

1. Introduction: the residue, and why not to attack it head-on
 2. Inherited data and the conservation-first stance
 3. Separability of the conserved normalization from BCB
 4. The Diagonal-Torus Lemma
 5. The Reversible-Mixing Dichotomy
 6. Why the inherited $U(1)$ holonomy is insufficient
 7. The two branches, and what settles each
 8. Relation to PAMV Open Problem 10 and the ODG-reduction paper
 9. Limitations and Open Problems
 10. Conclusion
-

Scope and Conditional Status

This paper continues the conditional-reconstruction tradition of PAMV, the UMP paper, the ODG paper, the Born-rule papers (Double Square Rule; Physical Necessity), and the ODG-reduction paper (the *Packing* paper). It does not claim to close the squaring residue. It claims to *reduce* the residue to a single substrate-level question and to prove that reduction by a sub-Hilbertian argument — one that at no point assumes an inner product, a quadratic form, or "reversible = isometry of a complex inner product."

Epistemic labelling is maintained throughout:

- **proven** — follows from the stated inputs (the inherited BCB additivity and diagonal- $U(1)$ phase conservation, the path-sum carrier H_0 , and standard functional/Lie-theoretic facts) by an argument given or cited here;
- **conditional** — follows given a named, separately-open hypothesis;
- **conjectural** — asserted as plausible with structural support but no argument approaching proof.

The three numbered results of the Abstract are **proven** in the stated sense. The resolution of the residue itself — which branch of the dichotomy holds — is **open**; this paper does not assert either branch.

A note on inheritance. The argument depends on two facts imported from prior papers, and the whole edifice is only as sound as they are. (i) From the Born-rule papers and the *Packing* paper: the outcome amplitude is a complex vector $\psi = (c_1, \dots, c_d) \in \mathbb{C}^d$ generated as a path-sum, with a per-path $U(1)$ phase (Physical Necessity Appendix C; *Packing* §3.5.2). (ii) From BCB (Bit Conservation and Balance): the substrate conserves a quantity that is additive over operationally disjoint outcomes. We use (i) and (ii) as inputs; we do not re-derive them. Following programme convention, time is emergent throughout — we speak of substrate ticks, commitment events, and reversible transport between configurations, never of motion through a pre-given temporal continuum.

1. Introduction: the residue, and why not to attack it head-on

The *Packing* paper completed a long arc. Where the question had been "does Operational Distinguishability Geometry descend from a substrate primitive?", that paper answered: four of six ODG components descend from finite packing plus equivariance; the carrier's *linear* structure and *complex phase* descend constructively from the path-sum substrate H_0 ; and exactly two things do not — the *refinement* completeness (Obstruction B) and the *squaring* of the inner product (the residue of Obstruction A). The squaring was shown to be no new gap but an inherited one: PAMV Open Problem 10, the question of whether the operational-volume/squared-norm identification $\text{Vol}_{\text{op}}(\chi) = \|\chi\|_{\mathbb{C}}^2$ is theorem or convention.

That paper also did the work of localizing the residue to its sharpest form and recording, in its §11.1, the right way to attack it. We take up exactly that recommendation. The residue, in the prior paper's three interderivable faces, is:

- (1) the reversible normalization-preserving symmetry group on \mathbb{C}^d is $U(d)$, preserving a positive-definite complex inner product;
- (2) the pairwise correlation kernel is positive-semidefinite;
- (3) selection is bilinear in paths.

The temptation is to attack face (1) affirmatively: assume the reversible dynamics, show it must be $U(d)$. The *Packing* paper diagnosed precisely why this fails — a unitary *is* a transformation preserving a positive-definite Hermitian inner product, so "the reversible group is $U(d)$ " presupposes the inner product whose origin is the question. Every affirmative route through face (1) re-imports the conclusion via the axiom "reversible = isometry of the distinguishability metric." Restating the residue as a theorem about the symmetry group, and then defining the

symmetry group as the isometries of the very norm in question, is closure-by-renaming. The prior paper refused it; so do we.

The alternative the prior paper recommended is **conservation-first**. Do not begin with the allowed transformations and ask what norm they preserve. Begin with the *conserved quantity* — fixed by bit conservation, which is a substrate law independent of any norm — and ask what *functional* it can be, and then ask what continuous transformations preserve *it*. The order matters: starting from the conserved functional and deriving its symmetry group cannot smuggle in $U(d)$, because $U(d)$ is never assumed; whereas starting from "reversible = $U(d)$ " assumes it on line one.

This paper carries out the conservation-first programme. The outcome is not a proof that the squaring is forced — we will be explicit that it is not, and that it *should not* be forcible from conservation alone. The outcome is a dichotomy: the squaring is forced **iff** one concrete, sub-Hilbertian, physically interpretable condition on the substrate holds, and is *not* forced otherwise. Naming that condition exactly — and proving the iff without ever assuming an inner product — is the contribution.

1.1 The shape of the result

The argument has three moves.

Move 1 (conservation \Rightarrow separability). BCB makes the conserved normalization additive over disjoint outcomes; the inherited per-path $U(1)$ phase makes it phase-blind. Together these force, in any fixed outcome basis, the separable form

$$N(\psi) = \sum_{i=1}^d h(|c_i|)$$

for a single function h . No exponent is selected: $h(r) = r^p$ gives the (p -th power of the) ℓ^p norm for every p , and all are admissible at this stage.

Move 2 (the Lie-theoretic dichotomy). A separable N has a large continuous symmetry group only in the ℓ^2 case. Precisely: the connected component of the group preserving N is the diagonal phase torus $U(1)^d$ for every separable N whose level sets are not Euclidean spheres, and enlarges to $U(d)$ exactly when N is the ℓ^2 norm. So a *continuous off-diagonal* symmetry — one mixing distinct outcome amplitudes — exists iff $p = 2$.

Move 3 (back to the substrate). The reversible dynamics of the substrate *is* a group of N -preserving transformations (it preserves the conserved bookkeeping quantity, by definition of conserved). Therefore: the dynamics contains a continuous off-diagonal generator iff the conserved N is ℓ^2 . Reading the iff in the useful direction — the substrate question "is there continuous reversible inter-outcome mixing?" decides the exponent.

The remainder makes each move precise (§§3–5), shows why the inherited $U(1)$ holonomy cannot by itself supply the off-diagonal generator (§6), lays out the two branches and what would settle each (§7), and connects the result to PAMV OP10 and the *Packing* paper (§8).

2. Inherited data and the conservation-first stance

2.1 The carrier (inherited, not re-derived)

From the Born-rule papers and *Packing* §3.5.2 (Input H0), the outcome amplitude of a macro-configuration with d operationally distinguishable outcomes is a complex vector

$$\psi = (c_1, \dots, c_d) \in \mathbb{C}^d, c_i = \sum_{P \in R_i} e^{i\theta(P)},$$

where R_i is the reversible-path set resolving outcome i and $\theta(P)$ is the inherited per-path $U(1)$ phase. Two structural facts come with this and are used as inputs:

- **(C1) Linearity over disjoint outcomes.** For disjoint outcome regions, $R_{\{i \cup j\}} = R_i \sqcup R_j$, so amplitudes add: the carrier is the free complex linear span of path contributions (H0).
- **(C2) Diagonal $U(1)$ phase.** Each coordinate carries an independent inherited phase freedom $c_i \mapsto e^{i\phi_i} c_i$ that is reversible and substrate-admissible (it is the per-path holonomy acting coherently within outcome i). The diagonal phase torus

$$\mathbb{T}_d := \{ \text{diag}(e^{i\phi_1}, \dots, e^{i\phi_d}) : \phi_i \in \mathbb{R} \} \cong U(1)^d$$

is therefore a subgroup of the reversible dynamics. **This is the part of the symmetry the prior papers genuinely established** — and, as §6 shows, it is exactly the part that does *not* select the exponent.

We emphasize what is *not* inherited: nothing in C1–C2 supplies any transformation outside \mathbb{T}_d (no continuous mixing between distinct coordinates), and nothing supplies a norm. The carrier is a complex vector space with a distinguished outcome basis and a diagonal phase symmetry. It is *not* yet a Hilbert space.

2.2 The conserved normalization (from BCB)

BCB asserts that the substrate conserves a bookkeeping quantity — the "bit content" balanced across commitment events. We extract the one feature we need and label it as an inherited axiom rather than re-deriving it:

Input BCB-N (Conserved normalization). There is a functional $N : \mathbb{C}^d \rightarrow \mathbb{R}_{\geq 0}$, the conserved normalization, with:

- **(N1) Disjoint additivity.** If outcome region M decomposes into operationally disjoint sub-outcomes with amplitude vectors restricted to disjoint coordinate blocks, N is additive across the blocks: $N(\psi \oplus \varphi) = N(\psi) + N(\varphi)$ for amplitudes supported on disjoint coordinate sets. (This is BCB: distinguishable outcomes carry independent, additive bit content.)
- **(N2) Reversible invariance.** N is invariant under every substrate-admissible reversible transformation U : $N(U\psi) = N(\psi)$. (This is what "conserved" means.)
- **(N3) Phase blindness.** N is invariant under \mathbb{T}_d : $N(\text{diag}(e^{i\varphi})\psi) = N(\psi)$. (This is N2 restricted to the inherited C2 phases, recorded separately because it is what we can use immediately.)
- **(N4) Faithfulness.** $N(\psi) = 0$ iff $\psi = 0$, and N is continuous.

We make no quadraticity, homogeneity-degree, or polarization assumption. N4's continuity is mild (the conserved quantity varies continuously with the configuration). N1 is the substantive BCB content; N3 is inherited; N2 is the definition of conserved and is where the substrate question will enter.

2.3 The stance

We will derive the *form* of N from N1, N3, N4 (Move 1, §3), establish the dichotomy purely from N 's symmetry (Move 2, §4), and only then re-engage N2 — the full reversible invariance — as the deciding substrate condition (Move 3, §5). The discipline is that $U(d)$ is never written down as an assumption; it appears, if at all, only as a *consequence* of an off-diagonal continuous generator existing, which is itself the open substrate question.

3. Separability of the conserved normalization from BCB

Proposition 3.1 (Separable form — proven, given BCB-N). Under N1 (disjoint additivity), N3 (phase blindness), and N4 (continuity, faithfulness), there is a continuous function $h : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ with $h(0) = 0$, $h(r) > 0$ for $r > 0$, such that in the distinguished outcome basis

$$N(\psi) = \sum_{i=1}^d h(|c_i|).$$

No exponent is selected: h is otherwise unconstrained at this stage.

Proof. Apply N1 to the decomposition of $\psi = (c_1, \dots, c_d)$ into its d single-coordinate amplitudes $\psi = e_1 c_1 \oplus \dots \oplus e_d c_d$ (the coordinate axes are operationally disjoint outcome blocks, since distinct outcomes are mutually distinguishable). Disjoint additivity gives

$$N(\psi) = \sum_i N(e_i c_i) = \sum_i n_i(c_i),$$

where $n_i(c) := N(e_i; c)$ is the single-coordinate normalization for outcome i . By N3, $n_i(c) = n_i(e^{\{i\}}c)$ for all φ , so n_i depends only on $|c|$: $n_i(c) = h_i(|c|)$. Outcome relabelling is a discrete substrate symmetry (distinguishable outcomes have no intrinsic order), so by N2 restricted to coordinate permutations the h_i are independent of i : $h_i = h$. Continuity and faithfulness (N4) give h continuous with $h(0) = 0$, $h(r) > 0$ for $r > 0$. ■

Remark 3.1.1 (Why the exponent is correctly *not* selected here). It is essential that Proposition 3.1 leaves h free. Were conservation plus phase alone to force $h(r) = r^2$, the squaring would be a triviality and the century of difficulty would be inexplicable. The separable form is the honest output of conservation-first: every ℓ^p normalization

$$N_p(\psi) = \sum_i |c_i|^p (h(r) = r^p), \quad 1 \leq p < \infty,$$

satisfies N1, N3, N4, as do many non-power functionals (any continuous h with $h(0)=0$). The exponent is undetermined by conservation and phase. What remains undetermined is precisely what the *off-diagonal* part of N2 — not yet used — will pin down. The squaring, if it comes, comes from the *symmetry* of N , not from conservation of N .

Remark 3.1.2 (Homogeneity is a separate, mild assumption — flagged). If one additionally posits that N is homogeneous of some degree, $N(\lambda\psi) = |\lambda|^k N(\psi)$, then $h(r) = r^k$ and $N = N_p$ with $p = k$, collapsing the function-space to the one-parameter ℓ^p family. We **do not** require homogeneity for the dichotomy of §§4–5 (the Diagonal-Torus Lemma holds for general separable N), but it is worth recording that the natural scaling of bit content (doubling all amplitudes multiplies content by a fixed factor) would supply it, reducing the residual freedom from "all continuous h " to "the exponent p ." We keep this as a **conditional** sharpening, not a needed step.

Remark 3.1.3 (The mixing forces an inner product without leaning on relabelling — what it forces, precisely). The proof of Proposition 3.1 used outcome-relabelling symmetry (N2 restricted to permutations) to set $h_i = h$. That is a real assumption, and one we would rather not lean on more than necessary. In the **Quantum branch** it is largely unnecessary, but the precise content must be stated carefully. Run the in-plane argument of §4.2 *without* assuming a common h — keep distinct h_k, h_ℓ on the connected pair. Then (★★) becomes

$$\beta h'_k(r_k) r_\ell \cos(\gamma + \Delta) + \beta' h'_\ell(r_\ell) r_k \cos(\gamma' - \Delta) = 0,$$

which forces $h'_k(r_k)/r_k = (\pm\beta'/\beta) \cdot h'_\ell(r_\ell)/r_\ell = \text{const}$, hence

$$h_k(r) = a_k r^2 \text{ and } h_\ell(r) = a_\ell r^2,$$

both quadratic — but with coefficients related only through the ratio β'/β , **not** forced equal. And $a_k \neq a_\ell$ is genuinely realizable: rescaling $u = \sqrt{a_k} c_k$, $v = \sqrt{a_\ell} c_\ell$ turns the off-diagonal generator into an ordinary $SO(2)$ rotation in (u, v) , which has both off-diagonal entries nonzero and integrates to an N -preserving flow for *any* $a_k, a_\ell > 0$. So what the mixing forces is "**N restricted to the connected pair is a positive-definite quadratic form**" — i.e. an inner product on the connected component — and *not* "an equal-weight ℓ^2 ."

This is exactly what the squaring requires, and the headline conclusion is unaffected (arguably strengthened): any positive quadratic form $a_k|c_k|^2 + a_\ell|c_\ell|^2$ is an inner product — absorb the weights into the basis normalization — so $p = 2$ and "the inner product is recovered" hold regardless of the coefficients. The genuine content of the remark survives: the **squaring needs no relabelling input** — quadraticity (hence the inner product, hence $p = 2$) is forced by the mixing alone, on each connected component, with off-diagonal connectivity propagating the quadratic form across all coordinates. What relabelling symmetry adds is only *equality of the coefficients* a_k — the isotropy of the inner product — which is either supplied by permutation symmetry or absorbed into basis normalization and is immaterial to $p = 2$. We record: the mixing derives the inner product without relabelling; relabelling (or basis normalization) supplies only its isotropy. The permutation assumption does genuine work only in the **Torus branch**, where it keeps the non- ℓ^2 conserved N symmetric across outcomes (a natural but not forced restriction; a Torus-branch model with coordinate-dependent h_i is conceivable and would carry an asymmetric non-quadratic normalization).

4. The Diagonal-Torus Lemma

This is the load-bearing mathematics, and it is sub-Hilbertian: it assumes only the separable functional of §3 and asks what *continuous* transformations preserve it. No inner product is assumed; the inner product, where it appears, is an *output*.

4.1 Statement

Lemma 4.1 (Diagonal-Torus Lemma — proven, with the C^1 upgrade H-reg below). Let $N(\psi) = \sum_i h(|c_i|)$ be a separable, phase-blind, faithful functional on \mathbb{C}^d , $d \geq 2$, with h continuous (Proposition 3.1) and **(H-reg)** $h \in C^1$ on $(0, \infty)$ with a one-sided derivative $h'(0^+)$ at the origin. Let $G_N \subseteq GL(d, \mathbb{C})$ be its invariance group, $G_N = \{ U : N(U\psi) = N(\psi) \ \forall \psi \}$, with identity component G_N^0 . Then:

- **(a)** $\mathbb{T}_d \subseteq G_N^0$ always (the diagonal phase torus preserves every separable phase-blind N).
- **(b)** If N is, up to a positive constant and a fixed power, the ℓ^2 norm — i.e. $h(r) = a \cdot r^2$ for some $a > 0$ — then $G_N^0 = U(d)$ (when N is taken as the squared norm; the isometry group of the norm itself is $U(d)$ likewise).
- **(c)** If N is *not* of that form — $h(r) \neq a \cdot r^2$ — then $G_N^0 = \mathbb{T}_d$. In particular G_N^0 contains **no** generator outside the diagonal torus: no continuous one-parameter subgroup of G_N mixes the amplitudes of two distinct coordinates.

Consequently: $G_N^0 \supseteq \mathbb{T}_d \Leftrightarrow N$ is the (scaled, fixed-power) ℓ^2 functional $\Leftrightarrow p = 2$. A continuous off-diagonal N -preserving symmetry exists if and only if N is Euclidean.

On H-reg. Proposition 3.1 delivers only continuity of h , so writing the infinitesimal condition (★) below presupposes a derivative we have not yet established; H-reg supplies it explicitly

rather than silently. It is mild and substrate-motivated: a conserved quantity that varies along smooth one-parameter reversible flows (which is what the off-diagonal generator we are testing *is*) is differentiable in the flow parameter, and separability transfers that to h . We carry H-reg as a named C^1 -upgrade hypothesis rather than fold it into N4, and flag that the lemma's conclusion ($h(r) = ar^2$) is itself smooth — so H-reg is consistent with the output and removable a posteriori, but it is not free a priori. **Status of the lemma: proven given H-reg; the C^1 -upgrade is conditional in the same sense as the phase-side continuity upgrade of the prior papers.**

4.2 Proof

(a) For $T = \text{diag}(e^{i\varphi_1}, \dots, e^{i\varphi_d})$, $|(T\psi)_i| = |e^{i\varphi_i} c_i| = |c_i|$, so $N(T\psi) = \sum h(|c_i|) = N(\psi)$. Thus $\mathbb{T}_d \subseteq G_N$, and \mathbb{T}_d is connected, so $\mathbb{T}_d \subseteq G_N^0$.

(b) For $h(r) = ar^2$, $N(\psi) = a \sum |c_i|^2 = a \|\psi\|_2^2$. The transformations preserving the standard Hermitian norm on \mathbb{C}^d are exactly $U(d)$, which is connected, so $G_N^0 = U(d)$. (That $U(d)$ preserves it is immediate; that nothing else does is the standard characterization of the unitary group as the full isometry group of the Euclidean Hermitian norm.)

(c) This is the substantive direction. Suppose G_N^0 contains a one-parameter subgroup $\exp(tX)$, $t \in \mathbb{R}$, with generator $X \in \mathfrak{gl}(d, \mathbb{C})$ **not** in the Lie algebra $\mathfrak{t}_d = \{ \text{diag}(i\varphi_1, \dots, i\varphi_d) \}$ of the diagonal torus. We show N must be Euclidean. Note " $X \notin \mathfrak{t}_d$ " does not immediately mean " X has an off-diagonal entry": X could be diagonal with a non-skew (non-imaginary) entry, e.g. $\text{diag}(1, 0, \dots, 0)$. We first exclude that by faithfulness, then handle the genuinely off-diagonal case.

Invariance $N(\exp(tX)\psi) = N(\psi)$ for all t, ψ , differentiated at $t = 0$ (using H-reg), gives the infinitesimal condition

$$\sum_i h'(|c_i|) \cdot (1/|c_i|) \cdot \text{Re}(\bar{c}_i (X\psi)_i) = 0 \text{ for all } \psi \text{ (with } c_i \neq 0), (\star)$$

where $(X\psi)_i = \sum_j X_{ij} c_j$.

Diagonal part is forced skew. Take ψ generic (all $c_i \neq 0$) and extract the part of (\star) coming from the diagonal entries $X_{ii} = s_i + it_i$: the contribution of coordinate i is $h'(|c_i|)/|c_i| \cdot \text{Re}(\bar{c}_i X_{ii} c_i) = h'(|c_i|)|c_i| \cdot s_i$. The off-diagonal contributions depend on the relative phases $\alpha_j - \alpha_i$ and average to zero under independent variation of those phases, while the diagonal piece is phase-independent. Hence the diagonal piece must vanish on its own:

$$\sum_i h'(|c_i|) |c_i| s_i = 0 \text{ for all radii } |c_i| > 0.$$

Since $h'(r)r$ is not identically zero on $(0, \infty)$ (else h constant, contradicting faithfulness), varying the radii independently forces every $s_i = 0$, i.e. $\text{Re}(X_{ii}) = 0$ for all i . Thus $\mathfrak{g}_N \cap \{\mathbf{diagonal}\} = \mathfrak{t}_d$: any diagonal generator is skew, hence already in the torus. So $X \notin \mathfrak{t}_d$ forces a genuine off-diagonal entry, say $X_{k\ell} \neq 0$ for some $k \neq \ell$. This also retroactively justifies dropping the diagonal terms below: they are torus terms (a), now shown to be the only diagonal content.

Out-of-plane leakage ($d \geq 3$). Before restricting to the $(e_{\underline{k}}, e_{\underline{\ell}})$ plane we must account for coordinates $m \notin \{k, \ell\}$ into which X may push amplitude. Setting $c_{\underline{m}} = 0$ on the plane, the flow gives $|w_{\underline{m}}(t)| = |t| \cdot |(X\psi)_m| + o(t)$ with $(X\psi)_m = X\{mk\}c_{\underline{k}} + X\{m\ell\}c_{\underline{\ell}}$, so coordinate m contributes to N a term $h(|w_{\underline{m}}(t)|) = h(|t| \cdot |(X\psi)_{\underline{m}}|) + o(t)$. Its one-sided t -derivative at 0 is $h'(0^+) \cdot |(X\psi)_{\underline{m}}|$, with *opposite* signs for $t \rightarrow 0^+$ and $t \rightarrow 0^-$ (a kink, since $|w_{\underline{m}}(t)| \propto |t|$). For $t \mapsto N(\exp(tX)\psi)$ to be constant — in particular differentiable — these kinks must cancel, and being one-sided-derivative terms of fixed sign they cannot cancel against the smooth in-plane terms. Hence for every $m \notin \{k, \ell\}$:

either $h'(0^+) = 0$, or $X_{\{mk\}} = X_{\{m\ell\}} = 0$.

In **both** cases the in-plane functional equation is untouched: in the first, every leaked term carries the factor $h'(0^+) = 0$; in the second, there is no leakage. So we may legitimately restrict to the $(e_{\underline{k}}, e_{\underline{\ell}})$ plane and the argument below runs for all $d \geq 2$.

The in-plane equation. On the plane write $c_{\underline{k}} = r_{\underline{k}} e^{i\alpha_{\underline{k}}}$, $c_{\underline{\ell}} = r_{\underline{\ell}} e^{i\alpha_{\underline{\ell}}}$, and let $X_{\{k\ell\}} = \beta e^{i\gamma}$, $X_{\{\ell k\}} = \beta' e^{i\gamma'}$ (the diagonal entries $X_{\{kk\}}$, $X_{\{\ell\ell\}}$ are skew by the step above and contribute only torus terms). The off-diagonal entries contribute the cross terms

$$h'(r_{\underline{k}})/r_{\underline{k}} \cdot \operatorname{Re}(r_{\underline{k}} e^{-i\alpha_{\underline{k}}} \cdot \beta e^{i\gamma} r_{\underline{\ell}} e^{i\alpha_{\underline{\ell}}}) + h'(r_{\underline{\ell}})/r_{\underline{\ell}} \cdot \operatorname{Re}(r_{\underline{\ell}} e^{-i\alpha_{\underline{\ell}}} \cdot \beta' e^{i\gamma'} r_{\underline{k}} e^{i\alpha_{\underline{k}}}) = 0.$$

Simplifying,

$$\beta h'(r_{\underline{k}}) r_{\underline{\ell}} \cos(\gamma + \alpha_{\underline{\ell}} - \alpha_{\underline{k}}) + \beta' h'(r_{\underline{\ell}}) r_{\underline{k}} \cos(\gamma' + \alpha_{\underline{k}} - \alpha_{\underline{\ell}}) = 0 \quad (\star\star)$$

for **all** $r_{\underline{k}}, r_{\underline{\ell}} > 0$ and all phases $\alpha_{\underline{k}}, \alpha_{\underline{\ell}}$. Vary the relative phase $\Delta := \alpha_{\underline{\ell}} - \alpha_{\underline{k}}$ freely. The two cosines are $\cos(\gamma + \Delta)$ and $\cos(\gamma' - \Delta)$; these are linearly independent functions of Δ unless $\gamma' = -\gamma \pmod{\pi}$, and even then $(\star\star)$ must hold identically in Δ , which forces each coefficient structure to be consistent across all Δ . Matching the Δ -dependence in $(\star\star)$ for all $r_{\underline{k}}, r_{\underline{\ell}}$ requires

$$\beta h'(r_{\underline{k}}) r_{\underline{\ell}} = \pm \beta' h'(r_{\underline{\ell}}) r_{\underline{k}} \text{ identically in } r_{\underline{k}}, r_{\underline{\ell}},$$

with the sign and phase fixed by γ, γ' . We now show **both** $\beta, \beta' \neq 0$, by ruling out the one-sided case directly rather than appealing to any properness of N . Suppose $\beta' = 0, \beta \neq 0$. Then $(\star\star)$ reads

$$\beta h'(r_{\underline{k}}) r_{\underline{\ell}} \cos(\gamma + \Delta) = 0 \text{ for all } r_{\underline{k}}, r_{\underline{\ell}} > 0 \text{ and all } \Delta.$$

Pick Δ with $\cos(\gamma + \Delta) \neq 0$; then $h'(r_{\underline{k}}) = 0$ for all $r_{\underline{k}} > 0$, so h is constant on $(0, \infty)$, hence $h \equiv 0$ by $h(0) = 0$ — contradicting faithfulness ($h(r) > 0$ for $r > 0$). By symmetry $\beta = 0, \beta' \neq 0$ is likewise impossible. A purely one-sided off-diagonal generator therefore cannot preserve a faithful separable N , so $\beta, \beta' \neq 0$. Then

$$h'(r_{\underline{k}})/r_{\underline{k}} = (\pm\beta'/\beta) \cdot h'(r_{\underline{\ell}})/r_{\underline{\ell}} \text{ for all } r_{\underline{k}}, r_{\underline{\ell}} > 0.$$

A function of r_k alone equal to a constant times a function of r_ℓ alone, for all independent r_k, r_ℓ , forces both sides constant: $h'(r)/r = 2a$ for some constant a , i.e.

$$h'(r) = 2a r \implies h(r) = a r^2 + h(0) = a r^2,$$

using $h(0) = 0$. Faithfulness ($h > 0$ for $r > 0$) forces $a > 0$. Thus $N(\psi) = a \sum |c_i|^2$, the Euclidean case. Contrapositively, if $h(r) \neq a r^2$ then no off-diagonal generator exists and $G_N^0 = \mathbb{T}_d$. ■

4.3 What the lemma is, and is not

Remark 4.3.1 (It is the complex Lamperti/Banach phenomenon, made infinitesimal).

Lemma 4.1(c) is the Lie-algebra form of a classical fact: for $p \neq 2$, the linear isometries of complex ℓ^p_d are generated by coordinate phase rotations and coordinate permutations (the connected component is the torus \mathbb{T}_d ; the permutations are discrete). The Euclidean case $p = 2$ is the unique member of the ℓ^p family whose isometry group is *continuous and acts transitively on the sphere* — equivalently, whose connected isometry component is larger than the diagonal torus. We give the infinitesimal proof directly (§4.2) rather than invoke the global classification, because the infinitesimal version is exactly what couples to "is there a continuous reversible generator?" — the substrate question — without any appeal to completeness or to the global isometry group.

Remark 4.3.2 (No inner product was assumed). The proof of (c) used only the separable form of N (Proposition 3.1) and the existence of one off-diagonal continuous generator. It produced $h(r) = ar^2$ — and *hence* the Euclidean inner product — as a conclusion. This is the discharge of the prior paper's worry: the inner product is not presupposed and then "found"; it is forced by the conjunction (separability \wedge off-diagonal continuity), neither conjunct of which is an inner-product assumption. Separability is BCB; off-diagonal continuity is the open substrate condition. The Hermitian inner product is their joint consequence.

5. The Reversible-Mixing Dichotomy

Now re-engage N_2 in full — the conserved normalization is invariant under *all* substrate-admissible reversible transformations, not merely the inherited diagonal phases. The reversible dynamics is, by definition of "conserved," a subgroup $R \subseteq G_N$: every admissible reversible U preserves N .

Theorem 5.1 (Reversible-Mixing Dichotomy — proven, given BCB-N and H0). Let N be the conserved normalization (separable by Proposition 3.1) and $R \subseteq G_N$ the group of substrate-admissible reversible transformations, with $R \supseteq \mathbb{T}_d$ (the inherited diagonal phases, C2). Exactly one of the following holds.

- **(Quantum branch.)** R contains a continuous one-parameter subgroup with an off-diagonal generator (continuous coherent redistribution of amplitude between two distinct

admissible outcomes). Then, by Lemma 4.1(c), N is the ℓ^2 norm up to scale; the conserved normalization is the squared Hilbert norm; the inner product of ODG component (i) is recovered; and PAMV Open Problem 10 closes **positively** — $\text{Vol}_{\text{op}}(\chi) = \|\chi\|_{\mathbb{C}}^2$ is a theorem. The squaring closes on this *single* generator. Separately, on the *size of the dynamics*: since $N = \ell^2$, Lemma 4.1(b) gives $G_{\mathbb{N}^0} = U(d)$, so $R^0 \subseteq U(d)$; and the supplied generator, together with the torus \mathbb{T}_d , generates at least a $U(2)$ mixing block on the connected pair, so $R^0 \supseteq U(2)_{\{k\ell\}} \times U(1)^{\{d-2\}}$. Whether $R^0 = U(d)$ — the *full* unitary dynamics — holds **iff** the substrate supplies off-diagonal generators connecting every coordinate pair (a spanning mixing graph); one generator does not force it. We do not claim $R^0 = U(d)$; the squaring residue closes on one generator regardless of how much of $U(d)$ the dynamics fills.

- **(Torus branch.)** $R^0 = \mathbb{T}_d$: substrate-admissible reversible dynamics is exhausted by diagonal per-outcome phase rotations together with discrete outcome relabellings. Then N is undetermined among all separable functionals (every ℓ^p , and more, is preserved); the conserved normalization is **not** forced to be ℓ^2 ; the inner product is independent structure, not a substrate consequence; and PAMV Open Problem 10 settles **negatively** — the squared-norm identification is convention, and a non- ℓ^2 substrate-admissible normalization exists.

These exhaust the possibilities (R^0 either does or does not exceed \mathbb{T}_d), and by Lemma 4.1 they are mutually exclusive.

Proof. $R \subseteq G_{\mathbb{N}}$ (\mathbb{N}^2 : reversible \implies N -preserving). $R^0 \supseteq \mathbb{T}_d$ (C2). Either $R^0 = \mathbb{T}_d$ (Torus branch) or $R^0 \supsetneq \mathbb{T}_d$. In the latter case R^0 contains a generator outside \mathfrak{t}_d — an off-diagonal generator (the diagonal-non-skew case is excluded by the faithfulness step in the proof of Lemma 4.1(c)) — so Lemma 4.1(c) forces N Euclidean; Lemma 4.1(b) then gives $G_{\mathbb{N}^0} = U(d)$, whence $R^0 \subseteq U(d)$, and bracketing the supplied generator with \mathbb{T}_d yields $R^0 \supseteq U(2)$ on the connected pair (Quantum branch). Note $R^0 = G_{\mathbb{N}^0} = U(d)$ is *not* concluded: R^0 is bounded above by $U(d)$ but filled only to the span of the supplied generators. The branches are exclusive because Lemma 4.1 makes "off-diagonal generator exists" equivalent to " N is ℓ^2 ," and a fixed N cannot be simultaneously ℓ^2 and non- ℓ^2 . ■

The residue, relocated. Theorem 5.1 is the paper's content. It converts the abstract face-(1) condition — "the reversible group is $U(d)$ " — into a concrete, sub-Hilbertian, physically meaningful substrate question:

The Squaring Question. Does substrate-admissible reversible dynamics include a continuous generator that coherently redistributes amplitude between two distinct admissible outcomes — a "substrate beam-splitter"?

"Why ℓ^2 ?" is the Squaring Question, exactly, and nothing more. The squaring is forced iff the substrate supports continuous reversible inter-outcome mixing, and is genuinely optional iff it does not. We have not decided which; we have proven that this single physical condition decides it, without assuming an inner product.

Remark 5.1.1 (The dichotomy is the quantum–classical divide). The two branches are not arbitrary alternatives; they are, structurally, *classical* versus *quantum* reversible probability, and naming them so makes the stakes vivid. Consider the Torus branch in its cleanest form: $N = N_1 = \Sigma_i |c_i|$ with reversible dynamics generated by per-outcome phases and outcome permutations. Strip the inherited (and here non-dynamical) $U(1)$ phases and what remains is exactly *classical reversible probability* — a probability vector $(|c_1|, \dots, |c_d|)$ summing to a conserved total, evolving by permutations (the reversible classical maps). The defining classical signature is precisely the absence of continuous inter-outcome mixing: classical reversible dynamics relabels alternatives, it does not coherently rotate one into another. The Quantum branch is the presence of exactly that mixing — the off-diagonal generator, the substrate beam-splitter — which is what makes reality quantum rather than classical. So the Squaring Question may be restated without loss as:

Is the substrate quantum or classical? — equivalently, does it support continuous reversible interconversion of distinct outcomes (quantum), or only their relabelling and independent rephasing (classical)?

The squared norm is the quantum signature; the failure to square is the classical one. This is a reframing of Theorem 5.1, not a new claim, and it costs nothing: the same off-diagonal generator that Lemma 4.1 ties to ℓ^2 is the one that distinguishes quantum from classical reversible probability.

6. Why the inherited $U(1)$ holonomy is insufficient

A natural hope is that the per-path $U(1)$ phase already established by the Born-rule papers (the holonomy argument, Physical Necessity Appendix C) *is* the continuous symmetry that selects ℓ^2 . It is not, and seeing why sharpens the open question.

Proposition 6.1 (The inherited holonomy lives in the torus — proven). The inherited per-path $U(1)$ phase acts on the outcome-amplitude vector as a diagonal transformation $c_i \mapsto e^{i\phi_i} c_i$, i.e. inside \mathbb{T}_d . It therefore preserves *every* separable phase-blind N (Lemma 4.1(a)), and selects no exponent.

Proof. The holonomy phase attaches to each *path* P and aggregates within the outcome region R_i that the path resolves. A coherent global phase on the paths resolving outcome i multiplies c_i by a unit modulus factor and leaves c_j ($j \neq i$) untouched — this is the definition of the diagonal action C2. Diagonal transformations are in \mathbb{T}_d ; by Lemma 4.1(a) they preserve all separable N . ■

This is the precise reason the prior papers, despite supplying complex phase, did not supply the squaring. **$U(1)$ holonomy is a *diagonal* continuous symmetry; the exponent is selected only by *off-diagonal* continuous symmetry.** The complex phase the programme worked hard to

derive is necessary for interference *within* an outcome (the path-sum recombination), but it is the wrong kind of motion to select ℓ^2 — it commutes with every ℓ^p . What is needed is categorically different: continuous coherent transport *between* distinct outcomes.

Remark 6.1.1 (Diagonal interference vs. inter-outcome interference). Two notions of interference must be distinguished, and the programme has so far secured only the first.

- *Intra-outcome interference* (secured): paths resolving the **same** outcome i add with phases, ψ within R_i exhibiting cancellation. This is the path-sum, the diagonal-torus content, consistent with all p .
- *Inter-outcome interference* (open): a reversible process coherently maps amplitude from outcome k into outcome ℓ and back — a rotation in the (e_k, e_ℓ) plane — so that the two outcomes are not merely independently phased but continuously interconvertible. This is the off-diagonal generator. It is the physical content of "superposition between distinct measurement outcomes can be reversibly rotated," and it is what beam-splitters, Rabi oscillations, and unitary gates instantiate in quantum mechanics.

The Squaring Question is whether the VERSF substrate supports inter-outcome interference as a continuous reversible process, or only intra-outcome interference plus discrete outcome relabelling. The prior papers establish the latter; the former is undecided.

Remark 6.1.2 (The qubit as the minimal instance — illustration, no new result). The dichotomy is stated at substrate level and assumes no quantum formalism, but its content is clearest in the minimal non-vacuous case $d = 2$ — a single pair of distinguishable outcomes. This is not merely a convenient example: by §9.6 the $d = 1$ configuration is degenerate (no distinct outcomes to mix, the Squaring Question vacuous), so a single such pair is the *smallest* object on which the residue lives at all. The familiar qubit is exactly one such pair, and the diagonal/off-diagonal split of Lemma 4.1 is visible on it directly.

Write a normalized 2-outcome amplitude in the conventional parametrization

$$\psi = \cos(\theta/2)|0\rangle + e^{i\varphi} \sin(\theta/2)|1\rangle,$$

drawn here on the ℓ^2 -sphere — and note immediately that the *sphere itself* is a Quantum-branch object, a point we return to below. The two kinds of continuous reversible motion separate cleanly:

- **Diagonal (torus) motions** rotate the relative phase between the outcomes,

$$|0\rangle \mapsto e^{i\alpha}|0\rangle, |1\rangle \mapsto e^{-i\alpha}|1\rangle,$$

i.e. motion in φ at fixed θ . These are the inherited $C2$ phases (Proposition 6.1), the σ_z -type generator, and they redistribute *no* amplitude between the outcomes — $|c_0|, |c_1|$ are unchanged. They preserve every separable N , every ℓ^p (Lemma 4.1(a)).

- **Off-diagonal motions** continuously transfer amplitude between the outcomes,

$$|0\rangle \mapsto \cos(\beta/2)|0\rangle + \sin(\beta/2)|1\rangle,$$

i.e. motion in θ — the σ_x -, σ_y -type generators. These are the off-diagonal generators of Theorem 5.1, and by Lemma 4.1(c) they preserve N **only** when N is ℓ^2 .

So the question "why does a single qubit carry continuous interconversion of its two outcomes, rather than only independent phase rotation of each?" is precisely the Squaring Question at $d = 2$ — the irreducible minimal form of the residue.

Two cautions keep the illustration faithful to the theorem. First, on the *carrier* the relevant connected group delivered by Lemma 4.1(b) is $U(2)$, not $SU(2)$: the Bloch sphere is the *projective* picture (states are rays), and the $SU(2) \cong \text{Spin}(3)$ acting on it, with $SU(2)/\{\pm 1\} = SO(3)$ the sphere-rotation group, is $U(2)$ after the global phase that the ray-picture discards is quotiented out. Whether one names the off-diagonal structure $U(2)$ (carrier level, matching the theorem) or $SU(2)/SO(3)$ (after removing global phase) is projective bookkeeping; it does not touch the diagonal/off-diagonal split, which is what the dichotomy turns on. Second, and more important: on the **Torus branch there is no Bloch sphere to rotate**. The non- ℓ^2 conserved level set is not a sphere — for $N_1 = |c_0| + |c_1|$ it is an ℓ^1 "diamond," for N_4 an ℓ^4 superellipsoid — and the diagonal torus acts on *that* body. Sphericity of the conserved level set and the existence of off-diagonal symmetry are the *same fact* (Lemma 4.1): the Quantum branch does not produce a sphere with some rotations missing, it produces the sphere in the first place. The Torus branch retains the diagonal phase motions acting on a non-spherical body; the full rotational geometry of the qubit simply does not emerge. The familiar geometry of a qubit is, in this sense, a concrete manifestation of the Quantum branch — present if and only if the substrate supports continuous reversible redistribution of amplitude between the two outcomes.

7. The two branches, and what settles each

The dichotomy is falsifiable in both directions, and each branch is a citable result.

7.1 The Torus branch (counterexample / negative settlement)

To establish it, construct: a substrate-admissible reversible dynamics on a $d \geq 2$ outcome configuration whose only continuous generators are diagonal (\mathbb{T}_d), with inter-outcome transformations restricted to discrete relabellings, such that the construction is consistent with TPB, BCB, temporal extensibility, and the path-sum carrier H_0 . Equip it with $N = N_p$, $p \neq 2$ (e.g. $p = 1$, the ℓ^1 "total amplitude" normalization, or $p = 4$). If such a dynamics is substrate-admissible, then by Theorem 5.1 the squaring is **not** forced, PAMV OP10 settles negatively, and the inner product is proven independent input.

Plausibility — and a sharper statement of where things actually stand. This branch is not merely *a priori* live; on the inherited inputs alone it looks close to manifestly admissible, and we say so directly rather than overstate the difficulty. Take the complex carrier (H_0), the conserved

normalization $N = N_1 = \sum_i |c_i|$, and reversible dynamics generated by per-outcome phases (C2) and outcome permutations. Then: disjoint additivity (N1) holds because moduli add; reversibility holds because permutations are reversible; the carrier is H_0 's; and N_1 is preserved by both phases (Lemma 4.1(a)) and permutations. This is just classical reversible probability riding on the complex carrier, with the inherited $U(1)$ phases absorbing intra-outcome interference without difficulty. **By every input this paper has used, it is admissible** — and if it is admissible, Theorem 5.1 settles PAMV OP10 *negatively, now*, and §7.1 is far closer to done than "attempt a construction" suggests.

The only thing that could block it is a constraint from **TPB** (tick-per-bit) — or from the *unextracted content of BCB* — that forbids classical-type reversible dynamics on the complex carrier and forces the off-diagonal generator into the admissible repertoire. And here is the honest tension the paper must own: **TPB is invoked in this section as the principle that might decide the branch, but it is nowhere stated as an axiom and nowhere used in any argument above** — and BCB was, by §2.2's own admission, only partially mined (we took N_1 and set the rest aside). So the whole analysis (§§2–6) runs on BCB-additivity (N_1), the diagonal phase C2, and the Lie lemma; neither TPB nor any BCB-surplus enters. The genuine state of play is a fork the paper has not resolved:

- **(a)** *The Torus construction is essentially trivial* — the classical-probability-on- H_0 model above is admissible against the full inherited principles (complete BCB *and* TPB), not merely against N_1 — in which case OP10 is settled negatively modulo that compatibility check, and the residue is *not* an open frontier but a nearly-closed negative result awaiting it.
- **(b)** *TPB or BCB-surplus forbids the restricted repertoire* — some content of the inherited principles beyond N_1 rules out classical-type dynamics on the complex carrier and forces continuous inter-outcome mixing — in which case the forbidding *is the entire content of the Quantum branch*, and a precise statement of the relevant principle plus a proof that it excludes the diagonal-only model is exactly the paper that would close OP10 positively.

We do not currently know which, because we have not engaged TPB at all, nor mined BCB past N_1 , and so have not checked the classical model against either. This is the paper's principal unmet obligation (§9.7), and we flag it as such rather than leave §7.1 reading as though the Torus construction were merely difficult. It is not obviously difficult; it is obviously *available* on the inputs used so far, and only an un-engaged principle (TPB, or BCB beyond N_1) could remove it. The next paper's target is therefore not "construct a Torus model" — that may be a paragraph — but "state TPB precisely, finish mining BCB, and determine whether either admits or forbids classical reversible probability on the H_0 carrier." That determination *is* the residue.

7.2 The Quantum branch (derivation / positive settlement)

To establish it, prove: that TPB, BCB, temporal extensibility, and H_0 *force* the existence of a continuous off-diagonal reversible generator — i.e. that the substrate necessarily supports inter-outcome interference as a continuous reversible process. Then Theorem 5.1 gives $p = 2$, the inner product is recovered, and PAMV OP10 closes positively.

What such a proof would need. It must produce, from substrate principles and *without* assuming a norm, a reversible flow that continuously interconverts two distinct admissible outcomes. The most promising line, flagged as **conjectural**: temporal extensibility plus the reversibility of commitment events may force that any two outcomes reachable from a common prior configuration are connected by a continuous reversible path in amplitude space (not merely by a discrete relabelling), because a discrete-only inter-outcome map would violate the continuity of substrate evolution that the holonomy continuity upgrade (Physical Necessity §C.8) posits. If commitment-event reversibility is genuinely continuous — the same continuity the phase-side argument seeks — then inter-outcome transport inherits that continuity, supplying the off-diagonal generator. This ties the Squaring Question to the *phase-side continuity upgrade*: both ask whether substrate reversibility is continuous, but for different motions (diagonal phase vs. off-diagonal mixing). They are **distinct** — continuity of diagonal phase does not entail continuity of off-diagonal mixing — but they may share a root in the continuity of commitment-event reversibility.

An internal gap within this conjecture, flagged at its own level. The line above contains a leap that must not be glossed: *continuous reachability of one outcome from another does not by itself yield an off-diagonal N-preserving Lie generator*. Three things are conflated and need separating. (i) A continuous reversible *path* connecting two outcomes need not be N-preserving *along its length* — only its endpoints need lie in the admissible set. (ii) Even an N-preserving path need not be a one-parameter *subgroup*; reachability is a statement about the orbit, not about group structure. (iii) Passing from "the reachable set is a connected manifold" to "there is a one-parameter subgroup $\exp(tX)$ with off-diagonal X" requires a homogeneity or group-closure input — that the admissible reversible maps form a (Lie) group acting transitively enough that the connecting paths are generated by a Lie algebra. None of (i)–(iii) is supplied by reachability alone. So the conjectural Quantum-branch argument has *two* layers, not one: first that inter-outcome transport is continuous (the part tied to the phase-side upgrade), and second that this continuity is *group-generated* rather than merely a continuous orbit. The second layer is the genuinely missing input, and we flag it as **conjectural with an unnamed structural premise** — strictly weaker in status than the phase-side upgrade, which at least has a named roadmap. Closing the Quantum branch requires both layers; this paper supplies neither and does not pretend the first entails the second.

7.3 The asymmetry between the branches

The branches are not symmetric in difficulty or in what they teach.

- The **Torus branch** is settled by a *construction* — exhibit one model — and is the lower-effort, higher-certainty move. A single explicit non- ℓ^2 substrate-admissible reversible dynamics ends the matter. We recommend it as the first attack, exactly as the *Packing* paper §11.1 recommended for the residue generally.
- The **Quantum branch** is settled by a *necessity proof* — show no diagonal-only substrate is admissible — and is harder, because it must exclude all restricted repertoires. It is also the more consequential outcome, since it would close PAMV OP10 positively and make the squared-amplitude Born rule a theorem of the substrate.

A rational research order, corrected by §7.1: the Torus model is *already available* on the inputs used here, so "attempt the construction first" is nearly vacuous — the construction is a paragraph. The real first move is to **finish mining BCB and state TPB, then check the classical model against both** (§9.7). If they admit the model, the residue is settled negatively and the inner product is independent structure. If either provably excludes every diagonal-only repertoire, that exclusion *is* the Quantum-branch proof, and the squaring is derived. Either way the deciding work is compatibility against the *fully* engaged inherited principles, not construction.

8. Relation to PAMV Open Problem 10 and the ODG-reduction paper

8.1 The residue is the same; the framing is sharper

The *Packing* paper's Proposition 3.2 proved the squaring residue identical to PAMV OP10 (whether $\text{Vol}_{\text{op}}(\chi) = \|\chi\|^2_{\mathbb{C}}$ is theorem or convention), and recorded the residue's three faces. The present paper does not introduce a new gap; it re-expresses the same residue as the Squaring Question and proves the re-expression (Theorem 5.1). The contribution relative to the *Packing* paper is the *conversion of an abstract symmetry-group condition into a concrete substrate condition*, achieved without assuming an inner product:

face (1): "reversible group is $U(d)$ " \implies (Theorem 5.1) \implies "substrate supports continuous off-diagonal reversible mixing."

The right-hand side is decidable at substrate level and does not presuppose its own answer. This is strictly more useful than face (1) as a research target: face (1) invites closure-by-renaming (define reversible as isometry); the Squaring Question forbids it (the off-diagonal generator either is or is not substrate-admissible, a fact about TPB/BCB, not about how one defines "reversible").

8.2 What flows back to the *Packing* paper

Either branch updates the *Packing* paper's ledger:

- **Quantum branch** \implies the squaring residue closes, ODG component (i)'s inner product is recovered, and the orthogonal/unitary upgrades of components (ii)–(iv) (which the *Packing* paper held conditional on the squaring residue) all discharge at once. The *Packing* paper's Theorem 8.1 Part 2 becomes unconditional, leaving only Obstruction B (FP3) and Input H0 outstanding.
- **Torus branch** \implies the squaring residue is proven *irreducible*; the *Packing* paper's "conditional on the squaring residue" rows become "conditional on an input now known

to be independent," and the programme's Hilbert structure is established as genuine extra information rather than a packing consequence — itself a clarifying result.

In neither case is Input H0 touched: the path-sum carrier is upstream of the Squaring Question (it supplies the \mathbb{C}^d on which N lives) and remains the second substrate primitive flagged in *Packing* §11.7. The present analysis in fact *uses* H0 (the complex carrier) and BCB as joint inputs, consistent with the *Packing* paper's accounting of two substrate debts.

8.3 Corroboration from the operational-reconstruction literature

The dichotomy is independently confirmed — and partly anticipated — by the axiomatic reconstructions of quantum theory, and honesty requires placing it against them. Hardy's 2001 framework and the Masanes–Müller 2011 derivation both single out quantum theory from within the broader class of operational/generalized-probabilistic theories, and in both the decisive axiom is a *continuity of reversible dynamics* postulate: that any pure state can be reversibly and *continuously* taken to any other (Hardy's continuity axiom; Masanes–Müller's "continuous reversibility"). Classical probability theory fails exactly this axiom — its reversible maps are the discrete permutations of pure states — while quantum theory satisfies it via the continuous unitary group. That is, structurally, our torus-vs-off-diagonal split: the reconstruction literature's continuous-reversibility postulate is the presence of continuous inter-outcome mixing, and its absence is the Torus branch (classical reversible probability). A clarification keeps this from reading as a conflation of two different strengths: Hardy/Masanes–Müller postulate continuity *plus transitivity* of the *dynamics* on pure states (any pure state reversibly reachable from any other), whereas Lemma 4.1's trigger is the existence of a *single* off-diagonal generator. These connect here via our own machinery, not by definition, though the connection must be stated at the right strength. One off-diagonal generator forces $N = \ell^2$ (Lemma 4.1(c)), whence the *norm's* invariance group is $G_{N^0} = U(d)$ (Lemma 4.1(b)), which is transitive on the unit sphere. That is the *available* symmetry — the ceiling the dynamics may occupy — and obtaining it as an output of a single generator, where the reconstructions assume the sphere's symmetry, is the strengthening. The reconstructions' postulate is properly matched only when the *dynamics* R fills enough of $U(d)$ to act transitively, which (per Theorem 5.1) requires off-diagonal generators spanning all coordinate pairs, not one. So the honest statement is: a single generator yields the ℓ^2 norm and hence the transitive symmetry *group* $G_{N^0} = U(d)$ as the available structure, and matches the reconstructions' transitive-*dynamics* postulate exactly when the substrate's mixing graph is connected — which is the same all-pairs-connectivity condition that distinguishes $R^0 = U(d)$ from a single $U(2)$ block. Either way the squaring ($N = \ell^2$) closes on one generator; the transitivity of the dynamics is the further, connectivity-dependent statement. The reconstructions thereby corroborate Theorem 5.1 by a completely independent route — operational axiomatics rather than substrate Lie theory — and confirm that "continuous interconversion of distinct pure outcomes" is indeed the hinge on which quantum-vs-classical turns.

Two things should be said plainly about what this does to the novelty claim. First, it *tempers* it: that continuous reversibility is the quantum/classical hinge is **known**, and we do not claim to have discovered the hinge. Second, it leaves the genuine contribution intact, and it is worth

stating exactly what survives. (i) **The sub-Hilbertian derivation of the iff.** Hardy and Masanes–Müller work within an operational framework that already presupposes a state space with its convex/probabilistic structure; we derive the equivalence "off-diagonal generator $\Leftrightarrow \ell^2$ " from a separable conserved functional and a Lie-algebra argument *without* assuming an inner product, a convex state space, or the reconstruction axioms — the inner product is an output (Remark 4.3.2). (ii) **The substrate-level relocation, with the continuity axiom explicitly refused as an assumption.** The reconstructions *postulate* continuous reversibility as an axiom selecting quantum theory; we *decline to postulate it* and instead identify it as the open substrate question, to be settled by TPB/BCB rather than assumed. The difference is exactly the difference between "assume the axiom that gives quantum theory" and "exhibit the substrate condition under which that axiom holds." (iii) **The conservation-first order.** We reach the hinge from bit conservation forward, not from the symmetry group backward, which is what keeps the derivation free of the circularity the *Packing* paper diagnosed. So the correct placement is: the dichotomy reproduces, at substrate level and without the inner product, the quantum/classical hinge known from operational reconstructions — corroboration of correctness, with the marginal contribution being the sub-Hilbertian route and the refusal to assume the continuity postulate that those reconstructions take as given.

9. Limitations and Open Problems

9.1 The Squaring Question itself (the residue, now concrete)

The paper does not decide whether substrate-admissible reversible dynamics contains a continuous off-diagonal generator. This is the residue, relocated but open. We record a **weak prior**, unchanged from the *Packing* paper's §11.1 in spirit but now sharper in object: we expect the inherited structure (diagonal phase + discrete relabelling) to be *insufficient* to force the off-diagonal generator, so that the Quantum branch, if it holds, must hold by an argument from the *continuity of commitment-event reversibility* (§7.2) and not from anything already established. Equivalently, we expect the Torus branch to be the default that must be actively excluded.

9.2 Is the off-diagonal generator forced by continuity of reversibility?

The conjectural Quantum-branch line (§7.2): does the continuity of substrate reversibility — the same continuity sought by the phase-side holonomy upgrade — extend from diagonal phase motions to off-diagonal inter-outcome motions? The two are logically distinct (Proposition 6.1: diagonal continuity is secured and selects nothing). A proof would need to show that a reversible process connecting two distinct outcomes cannot be purely discrete — that admissible reversibility is connected in a sense that includes inter-outcome paths. This is the single most consequential open sub-question, and it is **conjectural** with a named target.

9.3 The Torus construction — likely a paragraph, not a programme

The recommended first attack (§7.1): exhibit an explicit reversible dynamics on $d \geq 2$ outcomes with $R^0 = \mathbb{T}_d$ and a non- ℓ^2 conserved normalization. We now state more honestly than the "attempt a construction" framing suggests: on the inputs this paper has actually used (BCB-additivity, the diagonal phase C2, the H0 carrier), the model *already exists* — classical reversible probability ($N_1 = \sum |c_i|$, permutations plus phases) on the complex carrier (§7.1). The construction is essentially a paragraph, not an open programme. What is genuinely unsettled is not whether the model can be written down but whether **TPB admits it**; see §9.7. So the residue's negative settlement is closer than the *Packing* paper's "attempt the construction" recommendation implied — it awaits a TPB-compatibility check, not a construction.

9.4 Homogeneity and the function-space residue

§3 leaves h free among continuous functions; §4's Lemma holds for general separable N , so the dichotomy does not need homogeneity. But the Torus branch then carries a larger ambiguity (all separable N , not just ℓ^p). If one adopts the mild homogeneity sharpening (Remark 3.1.2), the Torus branch's freedom collapses to the exponent p , which is cleaner to reason about and to construct counterexamples within. Whether substrate bit-content scaling genuinely supplies homogeneity is a minor open point, flagged **conditional**.

9.5 Dependence on the inherited inputs

The argument is only as sound as BCB- N (the additivity N_1) and H0 (the complex carrier). If BCB additivity over disjoint outcomes fails — if conserved bit content is *not* additive across distinguishable outcomes — Proposition 3.1's separability fails and the analysis must be redone with a non-separable N , for which the Diagonal-Torus Lemma does not directly apply. We take N_1 as established by BCB; a referee should verify that the BCB papers license disjoint additivity in the strong, coordinate-blockwise form N_1 uses. This is the load-bearing inheritance, exactly as the *Packing* paper's Proposition 3.2 was load-bearing on the OP10 citation.

9.6 Beyond two outcomes

Lemma 4.1's proof works in a single (e_k, e_ℓ) plane and lifts to all of \mathbb{C}^d because an off-diagonal generator anywhere forces N Euclidean on that plane, and separability — **in the single-h form of Proposition 3.1** — then propagates the form to all coordinates: one plane quadratic \Rightarrow the shared h is quadratic \Rightarrow every coordinate is quadratic. The propagation is the *permutation symmetry* (single h) doing global work, not separability in any weaker sense. This must be distinguished from the 3.1.3 route, which deliberately drops permutation symmetry: there, one generator forces quadraticity only on its connected component, and global ℓ^2 additionally requires the off-diagonal mixing graph to span all coordinate pairs (the same all-pairs-connectivity that distinguishes $R^0 = U(d)$ from a single $U(2)$ block, §5). So the two routes carry different propagation hypotheses — single- h propagates from one generator; the

permutation-free route needs graph-connectivity — and §9.6's clean one-generator propagation is the single-h statement, not free on the 3.1.3 route. We have assumed $d \geq 2$ throughout; the case $d = 1$ is degenerate (no distinct outcomes to mix, the Squaring Question is vacuous, and probability is trivially the single conserved scalar). The dichotomy is a statement about genuine multi-outcome configurations.

9.7 State TPB precisely — and finish mining BCB — the principal unmet obligations

The deepest limitation is that **TPB is never stated as an axiom or used in any argument here**. It is invoked in §7 as the principle that might decide the branch, but the entire derivation (§§2–6) runs on BCB-additivity, the diagonal phase, and the Lie lemma — TPB does no work. This is not a cosmetic gap: §7.1 shows the Torus model is admissible on every input the paper *does* use, so the branch is decided entirely by whatever the engaged inputs leave out.

That phrasing — "whatever TPB adds" — is itself slightly too quick, and the correction names a *second* un-engaged candidate. §2.2 is explicit that we extracted only N1 (disjoint additivity) from BCB, "the one feature we need." BCB was therefore only *partially* mined, and its unextracted content is a co-decider sitting right alongside TPB. The name "Bit Conservation **and Balance**" hints that BCB carries a balance condition beyond bare additivity; if it does, that surplus could exclude the classical model without TPB doing anything. §9.5 worries about N1 *failing*; it does not consider BCB carrying *more* than N1. So the honest fork has two un-engaged candidate deciders, not one:

- **TPB**, never stated or used, and
- **BCB-surplus** ($\text{BCB} \not\supseteq \text{N1}$), the balance content beyond additivity that §2.2 deliberately set aside.

The Torus model is admissible against $\{\text{N1}, \text{C2}, \text{H0}, \text{Lie}\}$; it must be checked against *both* the full BCB (not just N1) and TPB before the negative settlement is secure. Either could forbid classical reversible probability on the H0 carrier on its own. The next paper's target is correspondingly twofold and sharply defined — it is **not** a construction (§9.3): (i) finish mining BCB — state its balance content precisely and check whether BCB-surplus alone excludes the diagonal-only repertoire; and (ii) state TPB precisely as a substrate axiom and check the same. A "no" from both settles PAMV Open Problem 10 negatively (the inner product is independent structure); a "yes" from either is the positive derivation. We flag this as the controlling open problem, superseding §§9.1–9.3 in priority: those describe the branches; this is what decides between them — and it names both un-engaged principles rather than presuming, as the bare "TPB" framing did, that BCB has already been fully mined when §2.2 says it has not.

10. Conclusion

The *Packing* paper reduced the origin of quantum probability geometry to two isolated obstructions and one inherited residue — the squaring — and recommended attacking the residue conservation-first rather than through the symmetry group. This paper carries out that recommendation and reaches a clean dichotomy.

From bit conservation alone (plus the inherited diagonal phase), the conserved normalization is separable: a sum over outcomes of a fixed function of each amplitude's size, with **no exponent selected** — correctly, since conservation should not by itself force the squaring. From a sub-Hilbertian Lie-theoretic fact — the Diagonal-Torus Lemma, proven here infinitesimally without assuming any inner product — the exponent is forced to 2 **iff** the conserved functional admits a continuous symmetry that mixes distinct outcomes, and is otherwise free. Reading this through the conserved-equals-reversible-invariant identity yields the Reversible-Mixing Dichotomy: the squaring is derived iff the substrate supports continuous reversible amplitude-redistribution *between* distinct outcomes — a substrate beam-splitter — and is genuinely independent structure if substrate reversibility is exhausted by per-outcome phase rotation plus discrete relabelling.

The century-old question "why the square?" is therefore reduced, exactly and without smuggled assumptions, to one concrete physical question:

Does substrate-admissible reversible dynamics include continuous off-diagonal mixing of distinct outcomes?

The inherited $U(1)$ holonomy — the complex phase the programme worked to derive — is provably the *wrong* motion to answer it: it is diagonal, commutes with every ℓ^p , and selects nothing. What selects the exponent is categorically different and not yet established: continuous inter-outcome interference, as distinct from the intra-outcome interference of the path-sum.

We do not decide the Squaring Question. We prove it is the only thing left to decide, that it is sub-Hilbertian, and — restated in the form Remark 5.1.1 makes plain — that it is the question *is the substrate quantum or classical?* The squared norm is the quantum signature, the off-diagonal generator is the substrate beam-splitter, and their absence is classical reversible probability; this is the same quantum/classical hinge the operational reconstructions (Hardy; Masanes–Müller) reach by postulating continuous reversibility, here derived at substrate level without assuming it.

One candid correction to the prior paper's recommended first move. The *Packing* paper, and the earlier sections here, suggested attacking the residue by *constructing* a non- ℓ^2 Torus model. We now see that construction is essentially trivial — classical reversible probability on the complex carrier is admissible on every input this paper uses (§7.1) — so the construction is not the frontier. The frontier is the inherited content this paper invokes or sets aside but never engages: **TPB**, which is nowhere stated or used, and the **unextracted balance content of BCB**, since §2.2 took only additivity (N1) and set the rest aside. The Torus model is available unless one of these forbids it. So the residue resolves to a sharply posed twofold obligation for the next paper: *finish mining BCB and state TPB precisely as substrate axioms, then determine whether either excludes classical reversible probability on the H_0 carrier.* A "no" from both settles PAMV Open Problem 10 negatively (the inner product is independent structure, reality is classical-on-the-carrier at the level of reversible dynamics); a "yes" from either is the positive derivation (the

substrate beam-splitter is forced, the squared-amplitude Born rule becomes a substrate theorem, and OP10 closes). Either way, the question is no longer "why the square?" but "does the substrate's full conservation-and-tick structure make it quantum?" — a single, answerable, substrate-level question.

The squaring is no longer a fog. It is a beam-splitter, present or absent in the substrate according to whether its full inherited structure (BCB beyond additivity, and TPB) demands it, and the whole of "why ℓ^2 " turns on which.