Deriving Operational Distinguishability Geometry from Finite Packing Structure

From Bounded Resolution to the Geometry of Quantum Probability — A Partial Reconstruction with Two Isolated Obstructions

Keith Taylor VERSF Theoretical Physics Programme versf-eos.com

---

For the General Reader

Quantum mechanics computes the chance of an outcome by *squaring* a number called the amplitude — the famous Born rule, $P = |\psi|^2$. Why squaring, and not some other operation, has puzzled physicists for a century. A longer-running VERSF project has been trying to *derive* that rule rather than assume it, by starting from something more primitive than Hilbert space and showing the quantum machinery is forced. This paper takes the deepest version of that question: where does the geometric stage on which quantum probability lives — what we call Operational Distinguishability Geometry, ODG — come from in the first place?

The proposed starting point is **finite distinguishability packing**: the simple idea that in any finite region of reality, at any finite resolution, only finitely many genuinely different states can be told apart. Think of it as a bound on how many distinct "snapshots" fit in a box. The question is how much of the quantum stage that single counting principle can rebuild.

The honest answer is: most of it, but not all, and the gaps are the interesting part. Finite packing — with a few mild and standard assumptions — rebuilds four of the six pieces of ODG: how the space splits into symmetry sectors, how measurements partition it, how it transports states around, and the packing structure itself. The *carrier* (the vector space amplitudes live in, with its complex numbers) is supplied by a separate, constructive route from the programme's other papers: amplitudes are built up as sums over paths, each carrying a phase, and those phases are forced to be complex by an argument about how time advances through irreversible commitments.

Two things genuinely do **not** follow, and the paper's contribution is to pin them down exactly rather than paper over them:

- **The squaring itself.** Building the carrier does not explain why probability is the *square* of the amplitude's size. Every route to "square" turns out to assume that the reversible motions of the theory preserve a particular notion of length — and that assumption is essentially the answer smuggled in. We show this leftover is an old, already-known

problem (PAMV Open Problem 10), now stated more sharply: *must reversible commitment dynamics preserve the squared length?* If yes, the squaring is derived; if no, it is an irreducible extra ingredient of reality. Either answer is worth having.
- **The refinement structure.** Whether every legitimate way of sub-dividing a measurement is one that the packing-and-resolution picture can already see. This is a new, separate question, and it could fail independently of the squaring.

So the paper is not "finite packing explains quantum mechanics." It is a **map**: here is what the counting principle rebuilds, here is what the path-sum construction adds, and here are the two precise, independent places where something genuinely quantum still has to be put in by hand. Naming exactly where the assumptions remain — and proving they are isolated rather than spread diffusely through the theory — is the point.

Abstract

The preceding VERSF reconstruction papers established that the Born rule is uniquely determined once probability is required to be compatible with Operational Distinguishability Geometry (ODG) — the unified operational structure comprising six components: carrier geometry, sector decomposition, projection structure, admissible transport, refinement structure, and finite distinguishability packing. In that programme ODG was a *structural input*, not derived; whether ODG itself descends from a more primitive substrate-level structure was left as the deepest open problem.

The present paper takes up that problem with finite distinguishability packing as the candidate primitive — the requirement that every operational region admits only finitely many mutually distinguishable states at finite resolution. The result is deliberately *not* presented as a clean derivation. It is a **partial reconstruction with two isolated obstructions**, and the paper's contribution is precisely the isolation: showing which ODG components descend from packing under transparent structural inputs, and which require a hypothesis that is not (yet) reducible to packing.

The accounting is as follows. Four of the six components — sector decomposition (ii), projection structure (iii), admissible transport (iv), and packing itself (vi) — descend from finite packing plus substrate equivariance by reasonably direct arguments (§§4–6). The remaining two are the obstructions:

- **Obstruction A (the inner-product gap).** Recovering the carrier *inner product* (component i) from the packing *norm* requires more than a normed vector space: it requires that the conserved, reversible-transport-invariant normalization on the carrier be the squared Hilbert norm. We resolve this obstruction into the part that is now closed and the part that is not. Drawing on the programme's Born-rule papers (the Double Square Rule and Physical Necessity papers), the carrier's **linear** structure and its **complex phase** are supplied constructively — amplitudes as linear path-sums over U(1) phases, with U(1) phase forced by tick-per-bit time, relational distinguishability, and temporal

extensibility (a complete impossibility theorem for finite holonomy, plus a conditional continuity upgrade). What is *not* supplied is the **squaring**: that the reversible normalization-preserving symmetry group is the unitary group U(d) preserving a positive-definite complex inner product. Every route to that step presupposes it — it is handed over by the axiom "reversible = isometry of the distinguishability metric," not derived from sub-Hilbertian inputs. We give this residue its three equivalent faces (unitary-invariance of the 2-norm ≡ positive-semidefiniteness of the correlation kernel ≡ bilinearity of selection), note that the classical negative-type / Schoenberg condition is one downstream characterization of it, and that it coincides with PAMV Open Problem 10 (strengthening $\mathrm{Vol\_op}(\chi) = \|\chi\|^2\_\mathbb{C}$ from convention to theorem). Obstruction A is therefore inherited, sharply located, and partly closed — we do not claim to close the squaring.
- **Obstruction B (the refinement gap).** Recovering refinement structure (component v) requires identifying admissible decompositions with resolution-stable packing refinements (hypothesis FP3). We show FP3 does not follow from the packing axioms FP1–FP2, isolate exactly what it adds, and characterize what its failure would cost.

The central result is therefore stated honestly as a conditional:

ODG Reduction Theorem (informal). Finite distinguishability packing plus substrate equivariance determines ODG components (ii)–(iv) and (vi) outright; *together with a second structural primitive — the path-sum substrate H0, independent of the packing axioms* — it determines the linear-and-complex carrier of component (i) constructively (the phase modulo a conditional continuity upgrade); the *inner product* of component (i) follows iff Obstruction A's squaring residue resolves (≡ PAMV Open Problem 10); and component (v) iff Obstruction B (FP3) resolves. The obstructions are independent and individually isolable, and H0 is carried as a named input throughout rather than folded into the packing axioms.

If both obstructions resolve, the VERSF probability hierarchy deepens by one layer, from

ODG → OIP → ODG-compatibility → UMP → Born rule

to

finite packing (+ H0) → ODG → OIP → ODG-compatibility → UMP → Born rule,

tracing the Born rule to the finite packing structure of distinguishable operational states *together with the path-sum substrate H0*. If only Obstruction B resolves, the reconstruction is complete modulo the squaring residue — itself an inherited gap. If neither resolves, the paper still delivers a four-of-six reduction, a constructive carrier, and a precise localization of the residue. In all branches the contribution is the same in kind: a map of *where* substrate-determination of ODG succeeds and where it is blocked.

Contents

Scope and Conditional Status

This paper sits in the same conditional-reconstruction tradition as PAMV, the UMP paper, and the ODG paper, and incorporates the constructive results of the programme's Born-rule papers (Double Square Rule; Physical Necessity). It does not claim to derive ODG from absolute first principles. It claims to *reduce* ODG to finite packing structure modulo two explicitly stated obstructions, and to show that those obstructions are the only places the reduction is incomplete.

A note on epistemic labelling, maintained throughout. We distinguish:

- **proven** — follows from the stated axioms (FP1, FP2), the inherited results of OG/URHG/PAMV, and the constructive results of the Born-rule papers, by an argument given or cited here;
- **conditional** — follows given a named hypothesis (H1, H2, FP3) or a flagged but-incomplete sub-result (the holonomy continuity upgrade) that is itself not established here;
- **conjectural** — asserted as plausible with supporting structural observation but no argument approaching proof.

The four reconstructed components (ii), (iii), (iv), (vi) are **conditional on H1–H2** (the metric-and-linear-completion hypotheses, mild and shared) but **not** on the obstruction hypotheses. Within component (i): the linear carrier and its complex phase are reconstructed constructively from **Input H0 (the path-sum substrate, §3.5.2)** — a second structural primitive, independent of FP1–FP2 — with the phase **conditional** on the holonomy continuity upgrade; the inner product is **conditional on Obstruction A's squaring residue**. Component (v) is **conditional on FP3 ≡ Obstruction B**. H0 is carried explicitly through the ledger and theorems rather than absorbed into the packing axioms, since the carrier's linearity rests on it and not on finite packing; the honest scope of the reduction is therefore *finite packing together with H0*, and the relationship between the two primitives is itself a reduction question (§§3.5.2, 11.7). We never label the assembled "ODG Reduction Theorem" as proven; it is a conditional whose conditions are named and, in one case (the squaring), shown inherited from a prior open problem.

What this paper does **not** do: it does not derive the squaring residue or FP3 from FP1–FP2; it does not derive H0 from FP1–FP2 (nor either primitive from the other); it does not establish that the six ODG components exhaust operational structure (an inherited convention, ODG paper §10.4); and it does not address the deeper substrate-origin questions (why $\mathbb{Z}_7$, why K = 7), which remain as the prior papers left them.

---

1. Introduction

The preceding VERSF probability papers established a staged reconstruction of the Born rule. PAMV showed that the minimal operational-probability axioms force quadraticity within coherence sectors. The UMP paper showed that the same quadratic rule must apply across operationally commensurable sectors, unifying PAMV's three bridging conditions under the Universal Measure Principle. The ODG paper identified the deeper geometric structure behind UMP — Operational Distinguishability Geometry — and characterized the Born rule as the unique invariant measure compatible with ODG. The Born-rule papers (the Double Square Rule and Physical Necessity) then attacked the carrier itself: deriving complex amplitudes as linear path-sums over geometric phases, and arguing that phase, pairwise selection, and the quadratic exponent are forced by physical admissibility. Throughout, ODG was treated as a structural *input*. The deepest open problem left by the ODG paper (its Open Problem 1) was whether ODG itself descends from substrate-level structure.

This paper takes up that problem. The candidate primitive is **finite distinguishability packing**: that every operational region admits only finitely many mutually distinguishable states at finite resolution. The guiding inversion is:

A physical state space is not first given as a smooth geometric object on which distinguishability is later measured. Rather, distinguishability is finite and operational from the outset, and geometry is the structure required to organize finite distinguishability consistently.

The reconstruction target is the six-component ODG structure:

1. carrier and inner product;
2. coherence-sector decomposition;
3. orthogonal projection structure;
4. admissible transport;
5. refinement structure;
6. finite distinguishability packing.

Component (vi) is the primitive input. The question is which of (i)–(v) it forces.

1.1 Why this is framed as a reduction, not a derivation

It would be easy, and overreaching, to state a single "ODG Reconstruction Theorem" asserting that all six components are forced. We decline that framing for a reason internal to the result's structure: it is not true without qualification, and the places where it fails are exactly the places of interest.

Of the five components to be reconstructed, three (sector decomposition, projection structure, admissible transport) follow from finite packing plus substrate equivariance by arguments that, while leaning on mild completion hypotheses, do not require anything beyond the packing data and the inherited symmetry. The other two do not reduce cleanly:

- **(i) the carrier and inner product** splits into a part that is now closed and a part that is not. The carrier's *linear* structure and *complex phase* are supplied constructively by the Born-rule papers (path-sums over U(1) phases); but the *inner product* — that the conserved reversible-transport-invariant normalization is the squared Hilbert norm — is not, and the step that would supply it presupposes it. We show this residual gap is identical to PAMV Open Problem 10, so it is an inherited obstruction, not a new one, and we decline to bury it inside a hypothesis labelled "compatibility."
- **(v) the refinement structure** requires the identification of admissible decomposition with resolution-stable packing refinement (FP3). This does not follow from the packing axioms; it is an additional structural commitment, and the strongest such commitment in the paper.

A reduction-with-isolated-obstructions framing is *more* informative than an overclaimed derivation: a sharply localized obstruction is a research target, whereas a theorem with a buried gap is a liability. The contribution is the localization itself — the demonstration that ODG-from-packing succeeds on four components, succeeds *constructively* on the linear-and-complex carrier, is blocked on the inner product's squaring by an inherited problem, and is blocked on (v) by one new hypothesis whose content we make fully explicit.

1.2 The central claim

Claim. Given finite distinguishability packing (FP1), packing equivariance (FP2), and the mild completion hypotheses (H1–H2): ODG components (ii), (iii), (iv), (vi) are reconstructed. Within component (i), the linear carrier and complex phase are reconstructed constructively (the phase modulo the holonomy continuity upgrade); the inner product is reconstructed iff Obstruction A's squaring residue resolves, and that residue ≡ PAMV Open Problem 10. Component (v) is reconstructed iff Obstruction B resolves; Obstruction B ≡ hypothesis FP3, which is shown independent of FP1–FP2.

The remainder develops the primitive data (§2), the four components that go through (§§4–6 plus the carrier-metric part of §3), the inner-product obstruction localized and partly closed (§3.4–§3.7), the refinement obstruction (§7), the assembly as an honest conditional (§8), the Born-rule consequence (§9), a head-to-head comparison of the obstructions (§10), and open problems (§11).

1.3 What this paper actually proves

A possible misreading would take the paper to claim that finite packing alone derives Hilbert space. That is not the claim. The result is more precise.

Finite packing, by itself, supplies bounded distinguishability capacity. With metric regularity and linear completion (H1–H2) it supplies a normed operational carrier; with substrate equivariance, sector decomposition; with packing-additive partitions, projection structure; with packing-preserving automorphisms, admissible transport. The Born-rule companion papers then supply a *separate* constructive route to the linear-and-complex carrier: amplitudes arise as path-sums over U(1) phase contributions. The carrier therefore need not be obtained by embedding an arbitrary metric space into Hilbert space — its linear structure is generated directly by composition of path contributions, and its complex phase by U(1) holonomy.

What remains unresolved is narrower and sharper:

Why is the conserved reversible normalization the squared Hilbert norm? Equivalently — why is the reversible symmetry group on outcome amplitudes U(d), rather than the isometry group of some other normed complex carrier?

This is the **squaring residue**. It is not created here: it is the issue already present in PAMV Open Problem 10 — whether $\mathrm{Vol\_op}(\chi) = \|\chi\|^2\_\mathbb{C}$ can be derived rather than adopted as a convention.

The contribution is therefore not a complete derivation of ODG from finite packing. It is a **reduction theorem with isolated residue**:

finite packing $\Longrightarrow$ ODG except for squaring and refinement.

The squaring residue concerns the *Hilbert character* of the carrier; the refinement residue concerns whether admissible decompositions are exhausted by resolution-stable packing refinements. These are independent obstructions, and isolating them makes the paper falsifiable at the structural level. To **defeat** the reduction, one must exhibit either (1) a substrate-admissible reversible dynamics whose conserved outcome normalization is not ℓ^2; or (2) an admissible decomposition not visible as a resolution-stable packing refinement. To **complete** it, one must prove either (1) that TPB, BCB, U(1) path-sum composition, and reversible admissibility force U(d) as the outcome-amplitude symmetry group; or (2) that every admissible decomposition is packing-visible and resolution-stable. Thus the paper reduces the origin of ODG to two precise questions — that is the central result.

The lesson is that the obstruction is no longer vague. The old question — *can finite packing derive Hilbert space?* — is replaced by a sharper one: *can reversible commitment dynamics force the ℓ^2 norm?* Finite packing alone does not know about squaring; path-sum composition supplies linearity; U(1) holonomy supplies complex phase. What remains is the claim that reversible probability-preserving dynamics must preserve the squared norm. If that can be derived, PAMV Open Problem 10 closes and the inner-product component of ODG is recovered; if it cannot, Hilbert structure is an irreducible input rather than a consequence of finite packing. Either outcome is useful: a positive result completes the finite-packing route to ODG; a negative result proves exactly where quantum structure enters as independent information.

The paper's central achievement is not that it eliminates every assumption, but that it shows exactly which assumptions remain and proves that they are isolated rather than diffuse.

---

2. Primitive Data: Finite Packing Structure

2.1 Operational states and distinguishability

Let $\mathcal{X}$ denote the primitive set of operational state representatives. At this stage $\mathcal{X}$ is *not* assumed to be a Hilbert space, a vector space, or even a metric space — only a set on which operational distinguishability can be evaluated.

For x, y ∈ $\mathcal{X}$ we write

$x \perp_\Delta y$

to mean that x and y are distinguishable at operational resolution Δ_{op}. A finite subset $\Sigma \subset \mathcal{X}$ is an **operational distinguishability set** if its elements are pairwise distinguishable at resolution Δ_{op}:

$\Sigma = \{x_1, \dots, x_N\} \subset \mathcal{X}$, with $x_i \perp_\Delta x_j$ for all $i \neq j$.

2.2 The finite packing axiom

Axiom FP1 (Finite distinguishability packing). For every operational region $M \subset \mathcal{X}$, every admissible distinguishability set $\Sigma(M) \subset M$ is finite, with cardinality bounded by an operational volume function:

$|\Sigma(M)| \leq Vol_{op}(M) / \Delta_{op}^{d_{op}}$,

where $Vol_{op}(M)$ is the operational packing volume of M, Δ_{op} is the minimal operational resolution, and d_{op} is the effective operational packing dimension.

This is the primitive structural input. The essential point — and the one that creates Obstruction A — is that **Vol_op is initially a counting measure, not a Hilbert norm.** It is defined by how many mutually distinguishable states fit inside a region at finite resolution. Nothing in FP1 says this counting measure is quadratic, polarization-compatible, or even derived from a norm at all.

Recovering Hilbert structure from a counting measure is precisely the work, and precisely where it is hard.

2.3 Packing equivalence

Two regions M, N $\subset \mathcal{X}$ are **packing-equivalent** if they support the same distinguishability capacity at all admissible resolutions:

$M \sim_{pack} N \Leftrightarrow Vol_{op}(M) = Vol_{op}(N)$.

Two states x, y $\in \mathcal{X}$ are **locally packing-equivalent**, written $x \sim_{loc\text{-}pack} y$, if every sufficiently small operational neighbourhood of x has the same packing profile as the corresponding neighbourhood of y. This is the first primitive equivalence relation, and the reconstruction begins from it.

---

3. From Packing Volume to Carrier Geometry — the Inner-Product Obstruction, Localized and Partly Closed

This section reconstructs as much carrier geometry as packing genuinely supplies — a metric and a norm (§3.1–3.3) — and then localizes precisely where the reconstruction of the *inner product* stops. The localization is the point of the section: drawing on the programme's Born-rule papers, the carrier's *linear* structure and *complex phase* are now supplied constructively (§3.5), so the obstruction is no longer the whole inner product but specifically the **squaring** — the claim that the conserved reversible-transport-invariant normalization is the squared Hilbert norm (§3.6). The classical negative-type/Schoenberg condition is one characterization of that residue, not its primary framing (§3.5.1).

3.1 Operational displacement and packing length

Let $\gamma : [0, 1] \to \mathcal{X}$ be an admissible operational path between representatives x, y $\in \mathcal{X}$. Define the **packing length** of γ by subdividing into minimal operationally distinguishable increments and taking the refinement limit:

$L_{pack}(\gamma) = \lim_{\max|\Delta t_i|\to 0} \Sigma_i [Vol_{op}(M_{\gamma(t_i), \gamma(t_{i+1})})]^{1/d_{op}}$,

where $M_{\gamma(t_i), \gamma(t_{i+1})}$ is the smallest operational region resolving the displacement from $\gamma(t_i)$ to $\gamma(t_{i+1})$ within the ambient operational geometry — a d_{op}-dimensional resolving region (a ball), so Vol_{op} here is the **ambient region-capacity** and scales as $\varepsilon^{d_{op}}$. The exponent $1/d_{op}$ converts this d_{op}-dimensional packing volume into a linear displacement, so that L_{pack} scales as a length. (This ambient region-capacity is a distinct functional from the *two-state plane-capacity* $C(x, y)$ that enters the squared-norm identification of §3.2 and Proposition 3.2; both derive from the same FP1 counting data but are computed on different region types and carry different intrinsic dimensions. Remark 3.2.1 reconciles them.)

The induced **packing distance** is

$$d_{pack}(x, y) = \inf_{\gamma : x \to y} L_{pack}(\gamma).$$

3.2 The metric (mild hypothesis)

Hypothesis H1 (Metric regularity). d_{pack} is finite, symmetric, non-degenerate up to operational indistinguishability, and satisfies the triangle inequality on equivalence classes of operationally distinguishable states.

H1 is mild: symmetry and the triangle inequality are near-automatic for an infimum-over-paths construction, and non-degeneracy up to indistinguishability is definitional (states at zero packing distance are operationally identical). Under H1 the quotient $\mathcal{X} / \sim_{op}$ inherits a genuine metric. This is the first reconstructed object: a **metric operational carrier**.

3.3 The norm (mild hypothesis)

Suppose the carrier admits a distinguished null state 0 representing zero distinguishability content. Define the **packing norm**

$$\|x\|_{pack} := d_{pack}(x, 0).$$

Hypothesis H2 (Linear completion). The metric completion of finite operational superpositions of packing-distinguishable representatives admits a vector-space structure compatible with $\|\cdot\|_{pack}$.

H2 is also comparatively mild *in the VERSF setting*: the OG carrier-construction papers already motivate admissible superposition, so H2 imports an inherited structure rather than inventing one. Under H1–H2 the carrier becomes a normed vector space $\mathcal{A}_{pack}$.

Proposition 3.1 (Metric/norm carrier — conditional on H1–H2). Under FP1, FP2 (§4), and H1–H2, finite packing reconstructs a normed operational carrier ($\mathcal{A}_{pack}$, $\|\cdot\|_{pack}$). **Status: conditional on H1–H2.**

Proof. FP1 supplies finite operational volume and bounded distinguishability capacity. H1 promotes packing distinguishability to a metric on operational equivalence classes. H2 supplies the linear completion required for superposition and a norm compatible with the metric. ■

So far so good — and crucially, this much does *not* require the inner product. The carrier is normed, not yet inner-product. The next step is where it breaks.

3.4 Obstruction A: the inner product does not follow from the norm

A norm determines an inner product **iff** it satisfies the parallelogram identity

$\|x + y\|^2\_pack + \|x - y\|^2\_pack = 2\|x\|^2\_pack + 2\|y\|^2\_pack$,

in which case polarization recovers the inner product. One might try to assert the parallelogram identity as a "polarization-compatibility hypothesis" and proceed as though the inner product had thereby been *derived*. That move is circular. **Assuming the parallelogram identity is not a regularity condition on the way to an inner product; it is the inner product, restated.** A packing norm built from a counting measure has no a priori reason to be quadratic — generic norms (ℓ^p for $p \neq 2$, polytope gauges, sub-Finsler lengths) are not.

We therefore state the gap as an obstruction:

Obstruction A (Inner-product gap). Finite packing plus H1–H2 yields a normed carrier. It does *not* by itself yield the inner product. The reconstruction of ODG component (i) (carrier *and inner product*) is blocked at exactly this point. Recovering the inner product requires an independent argument that the conserved carrier normalization is quadratic.

Remark 3.4.0 (Why the *diagonal* packing norm cannot close Obstruction A — an arity argument, correctly scoped). It is worth stating why attempts to force the parallelogram identity from the packing *norm* have repeatedly relocated the assumption rather than discharged it. The derived norm $\|x\|\_pack = d\_pack(x, 0)$ is a **unary** functional: it scores an individual state by its packing distance from the null state. The parallelogram identity is a **binary** law on pairs x, y. One cannot derive the binary law from the unary diagonal norm without positing how unary scores aggregate *across pairs* — and that positing is where quadraticity re-enters in disguise. (Concretely: any "isotropy of plane-capacity C(x, y)" hypothesis that yields the parallelogram law must define C(x, y) as a sum of squared norms; on an ℓ^p carrier the natural plane-capacity aggregates as an ℓ^p quantity and isotropy gives the ℓ^p relation, not the Euclidean one. Quadraticity is conserved through such arguments, never created.)

The scope of this claim is exact and worth stating precisely. The argument forbids deriving the binary law *from the unary diagonal norm* — it does **not** forbid it from the full packing functional. Vol_op(M) evaluated on a two-state region $M_{\{x,y\}}$ is precisely the binary plane-capacity C(x, y); the full packing structure therefore does carry binary content, and it is only the *diagonal restriction* $\|x\|\_pack$ = (Vol_op of a region anchored at x) that is unary. So the correct conclusion is not "packing cannot, structurally, force quadraticity" but rather: *the diagonal norm is the wrong restriction to start from, because it discards the binary content that the parallelogram law constrains.* The live question for the full packing functional — does Vol_op on two-state regions force the negative-type / quadratic relation? — is exactly the question §3.5.1 takes up under the relational primitive D, of which the two-state plane-capacity is the natural realization. This is why the rest of §3 abandons the diagonal norm in favour of the relational primitive: not because packing is structurally barred from carrying the answer, but because the binary content lives in Vol_op-over-pairs, which *is* D, and that is where the question should be posed.

3.5 Two routes to the carrier, and what each leaves open

There are two non-question-begging routes past Obstruction A. The first (relational) sharpens the gap to a named classical condition; the second (constructive), supplied by the Born-rule papers, *closes the linear-and-complex part of the carrier outright* and relocates the residue to a single sharp condition. We take them in turn, then state the residue and the honest status.

3.5.1 The relational route (Schoenberg negative type) — one characterization of the residue

Take the **relational** distinguishability quantity as primitive rather than the unary packing volume.

Definition 3.3 (Relational distinguishability). Let $D : \mathcal{X} \times \mathcal{X} \to \mathbb{R}_{\geq 0}$ assign to each pair of operational states the minimum operational resource required to distinguish them, subject to symmetry (D1: $D(x, y) = D(y, x)$), identity (D2: $D(x, y) = 0$ iff $x \sim_{op} y$), a composition-consistency law (D3), refinement stability (D4), and distinguishability conservation (D5). D is a binary object from the outset — the right arity, per Remark 3.4.0 — and the packing volume is recovered downstream as the diagonal-anchored restriction $Vol_{op}(x) \propto D(x, 0)$, so nothing in §§4–7 is lost.

Schoenberg's theorem (the embedding criterion). A symmetric D with $D(x, x) = 0$ embeds isometrically into a Hilbert space — is realizable as a squared inner-product distance $D(x, y) = \|\Phi(x) - \Phi(y)\|^2_{\mathcal{H}}$ — **if and only if D is of negative type**: for every finite set $x_1, \ldots, x_n$ and real coefficients $c_1, \ldots, c_n$ with $\Sigma_i c_i = 0$,

$$\Sigma_i \Sigma_j c_i c_j D(x_i, x_j) \leq 0.$$

Negative type is exactly the condition separating ℓ^2 from ℓ^p ($p \neq 2$) and from all non-Hilbertian metric spaces. It is **not** a renaming of the parallelogram identity: it is an infinite family of inequalities over all finite point-configurations, with independent geometric meaning. So the relational route sharpens Obstruction A to: *do (D1)–(D5) force negative type, or is it the irreducible residue?* Symmetry, identity, and a triangle-type composition give a metric but **not** negative type (ℓ^p, $p > 2$, satisfies all three and fails it), so the weight falls on (D3)/(D4)/(D5).

Two cautions, both decisive for how this route is billed. First, **Schoenberg supplies a metric embedding, not a linear carrier**: a set may embed isometrically into $\mathcal{H}$ without its image being a subspace (a circle embeds in the negative-type sense; its image is not closed under +). So negative type alone does not give the vector-space carrier on which sectors, projections, and transport act. Second, and consequently, **the relational route is best read as one *characterization* of the inner-product residue, not its primary resolution** — it tells us *which* relations are Hilbert-realizable, but it neither supplies the linear carrier nor, on its own, the complex phase. Those come from the constructive route.

3.5.2 The constructive route (path-sums over U(1) phase) — closes the linear-and-complex carrier

The Born-rule papers supply the carrier directly rather than embedding a pre-existing set, and this closes the part of Obstruction A that the relational route leaves open.

Linear carrier (Double Square Rule §3.5, §6.1). The construction imports one structural input beyond the packing data, which we name explicitly rather than absorb silently:

Input H0 (Path-sum substrate). Reversible micro-paths exist, carry a U(1) phase, and compose by amplitude addition: for a macro-outcome A with reversible-path set R_A, the amplitude is $\psi_A = \Sigma_{P \in R_A} e^{i\theta(P)}$, with the per-path phase sourced from the holonomy argument below.

Given H0, the linear structure follows by construction. For disjoint outcome regions, $R_{A \cup B} = R_A \sqcup R_B$, whence $\psi_{A \cup B} = \psi_A + \psi_B$ — operational composition $\oplus$ **is** amplitude addition (union of disjoint path-sets), and scalar action is multiplication of phases by a constant. The carrier is the free linear span of path contributions, generated linearly rather than embedded. **This makes H2's linear-completion content constructive rather than merely hypothesized, and it leaves no residual "rotation of an embedded image" freedom to exclude** — so the embed-then-linearize route (recovering a linear carrier from a bare metric embedding via Mazur–Ulam isometry-affinity arguments) is not needed: linearity is supplied directly, not reconstructed from a metric set. **Status: the linear-carrier part of Obstruction A is closed (proven, given H0).**

On the relationship between H0 and finite packing. H0 is a *second structural primitive*, not a consequence of FP1–FP2: finite packing bounds distinguishability capacity but says nothing about reversible paths or phase composition, and the path-sum substrate says nothing about packing volume. They are presented here as independent imports. Whether both descend from a common substrate root (the TPB/BCB tick-and-commitment structure that the Born-rule papers and the packing axiom both invoke informally) is plausible within the VERSF programme but is **not established here**, and is flagged as a further reduction question (§11.7). Consequently the title's "from finite packing structure" is shorthand: the carrier specifically rests on H0, and the honest scope of the reduction is *finite packing together with the path-sum substrate*. We carry H0 explicitly through the ledger (§8.1), Theorem 8.1, and the Born-rule corollary (§9) so that no result is credited to finite packing alone that in fact leans on H0.

Complex phase (Physical Necessity §2, Appendix C). The construction needs the per-path factor $e^{i\theta(P)}$ to be genuinely U(1)-valued — else ψ_A is a real count with no interference. The holonomy argument forces continuous U(1) phase from tick-per-bit time (TPB), relational distinguishability, and temporal extensibility, with no Hilbert space assumed, in two layers of differing status:

- **Layer 1 (Appendix C.4): finite holonomy is impossible.** A finite holonomy set H of size K caps operationally distinguishable commitment histories at K, but TPB plus temporal extensibility force that number to grow without bound — contradiction for $n \geq K$. **This is a complete combinatorial proof — proven.**

- **Layer 2 (Appendix C.5–C.6): infinite → continuous → U(1).** Upgrading to a continuous one-parameter group uses incremental and total boundedness plus a Gleason–Montgomery–Zippin structure theorem, with four technical items flagged for full rigor (§C.8). **Conditional** — and, by the programme's own statement, the sole admitted open obligation on the phase side.

Status: the complex-phase part of Obstruction A rests on a chain — TPB → U(1) holonomy → complex per-path phase — whose first link is proven and whose continuity link is conditional, honestly flagged. This supersedes any separate "real → complex" sub-step routed through $\mathbb{Z}_7$-equivariance: the complex phase is the natural content of U(1) holonomy.

3.5.3 The relocated residue: the squaring

What §3.5.2 supplies is the amplitude *vector* $\psi = (c_1, \ldots, c_d) \in \mathbb{C}^d$. What it does **not** supply is that probabilities are $|c_i|^2$ — that the conserved normalization is the ℓ^2 norm. This is the genuine residue of Obstruction A, and the Born-rule papers locate it precisely.

Every route to the exponent $p = 2$ in those papers runs through one gate: invariance of the normalization under the **unitary group**. *Double Square* §5.6 excludes $p \neq 2$ because "$\|\psi\|^p$ is not preserved under unitary evolution"; *Physical Necessity* §8 makes $p = 2$ "the unique TPB-consistent exponent" because $\Sigma|c_i|^p$ is invariant under all unitaries iff $p = 2$, and §4 excludes higher-order kernels via the Mixed-Symmetric-Powers Lemma (a U(d)-invariant Hermitian form on the defining irrep $\mathbb{C}^d$ is, by Schur, unique up to a positive scalar — namely the standard form $\langle c, c\rangle$ — and $\|c\|_2^2$ is that form evaluated on the diagonal; higher even powers such as $\|c\|_2^4$ are of course also U(d)-invariant as *polynomials*, so the restriction doing the work is that the selection functional is a Hermitian/quadratic form, not the bare invariance); the *Entropy-Momentum* paper's four Born routes each presuppose the complex inner product outright (Gleason–Busch assumes Hilbert space; the Fubini–Study route writes the Born map $B(\psi) = (|\langle 0|\psi\rangle|^2, \ldots)$ into its setup; the martingale route takes $p_i(0) = |\alpha_i|^2$ as initial datum). But a unitary *is* a transformation preserving a positive-definite Hermitian inner product, so "normalization invariant under U(d)" presupposes the inner-product structure whose origin is the question. The 2-norm is singled out *because it is the norm the assumed symmetry group preserves* — Gleason's content, relocated rather than eliminated.

The residue, at maximum precision. The holonomy argument delivers a U(1) phase *per path*. The Born rule needs the full group U(d) acting on the *outcome-amplitude vector* $\mathbb{C}^d$ and preserving $\Sigma|c_i|^2$. The jump

per-path U(1) phase $\Longrightarrow$ U(d) on $\mathbb{C}^d$ preserving the ℓ^2 norm

is the inner-product-installation step, and it is **not derived**: it is handed over by the axiom that reversible maps are isometries of the pre-existing distinguishability metric (Double Square A4; Physical Necessity's "reversible = isometry"), after which §8 simply *demands* invariance of $\Sigma|c_i|^p$ under that group.

Three equivalent faces of the one residue — named together so they do not masquerade as separate gaps:

1. the reversible normalization-preserving symmetry group is U(d) preserving a positive-definite complex inner product;
2. the pairwise correlation kernel K is positive-semidefinite (Double Square §5.2–5.3, the PSD/Mercer step);
3. selection is bilinear in paths (Double Square Axiom A7, the "pairwise principle").

These are interderivable given the rest of the structure. **The negative-type condition of §3.5.1 is a fourth, downstream characterization of the same residue, one level below — a Hilbert-realizability test on the metric D rather than a condition on the symmetry group on outcomes.** The word "downstream" is earned by a one-directional implication, which we state explicitly. If the residue resolves — the reversible normalization-preserving group is U(d) on a positive-definite complex inner product — then the induced relational metric $D(x, y) = \|\Phi(x) - \Phi(y)\|^2_{\mathscr{H}}$ is, by construction, a squared Hilbert distance, hence of negative type (Schoenberg). So **symmetry-group/inner-product condition ⟹ negative type of the induced D.** The converse fails: negative type of D yields (Schoenberg) only a *metric embedding* of the state set into $\mathscr{H}$, not the linear carrier and not the unitary action on outcomes — a negative-type D can be realized on a set whose image is not a subspace and on which no U(d) action is implied (§3.5.1's caution). So negative type does **not** return the symmetry-group condition. The implication is therefore strictly one-way, condition (1) ⟹ negative type and ⇍, which is exactly what licenses "shadow": negative type is the trace the residue leaves on the relational metric, weaker than the residue itself. The correct primary home of "why Hilbert" is the symmetry-group condition (1); negative type is its one-way shadow.

3.6 Obstruction A, restated and located; relation to PAMV Open Problem 10

Obstruction A (restated). Within ODG component (i): the *linear carrier* and *complex phase* are reconstructed constructively (§3.5.2), the phase modulo the holonomy continuity upgrade. The *inner product* — the squaring — is **not** reconstructed; it is relocated to the U(d)-invariance / PSD-kernel / A7 condition (§3.5.3) and there assumed via "reversible = isometry," not derived from sub-Hilbertian inputs.

Proposition 3.2 (The squaring residue ≡ PAMV Open Problem 10). The squaring residue of Obstruction A is identical to PAMV Open Problem 10 — whether $\mathrm{Vol\_op}(\chi) = \|\chi\|^2_{\mathbb{C}}$ can be strengthened from definitional convention to derived theorem.

Argument. If $\mathrm{Vol\_op}(\chi) = \|\chi\|^2_{\mathbb{C}}$ holds as a theorem, the packing volume is by construction the square of an inner-product norm, the conserved normalization is the 2-norm, and the residue dissolves. Conversely, if the conserved carrier normalization is the squared Hilbert norm (the residue resolves), then on the diagonal-anchored restriction $\mathrm{Vol\_op}(x) \propto D(x, 0)$ the packing volume is a quadratic function of inner-product displacement, which is $\mathrm{Vol\_op}(\chi) = \|\chi\|^2_{\mathbb{C}}$ up to normalization. The two are inter-derivable. ■

Remark 3.2.1 (Reconciling the two roles of Vol_op — exponent bookkeeping). The inter-derivability above requires care, because Vol_op appears in two distinct functionals at different scalings, and "up to normalization" must not be allowed to absorb an exponent silently. The two functionals must be named, because the symbol Vol_op alone does not distinguish them:

- the **ambient region-capacity**, Vol_op of the smallest d_op-dimensional region (a ball) resolving a displacement, used in §3.1's L_pack as [Vol_op]^{1/d_op}; this scales as length^{d_op};
- the **two-state plane-capacity** C(x, y), Vol_op restricted to the 2-plane section spanned by x and y (equivalently, the binary realization of the relational D(x, y) of §3.5.1); this is intrinsically a 2-dimensional capacity (an area) and scales as length^2.

The reviewer's tension — that a finite two-state region M_{x,y} and an infinitesimal increment M_{γ(t_i),γ(t_{i+1})} look like the same kind of object, so assigning one dimension 2 and the other dimension d_op forces d_op = 2 — is resolved by noting these are **different restrictions of the same FP1 counting data, not the same functional on different-sized regions**. The increment in §3.1 is the *ambient* resolving region: it must resolve the displacement within the full local geometry, hence carries the ambient packing dimension d_op. The plane-capacity C(x, y) is a *sectional* restriction: the distinguishability capacity computed within the 2-plane through x and y, hence dimension 2 regardless of d_op. The same FP1 functional yields an ε^{d_op} ball-capacity on an ambient increment and an $\|\cdot\|^2$ area-capacity on a 2-plane section; no global d_op = 2 is implied.

Bridging to OP10. PAMV Open Problem 10 is stated for the FP1 object Vol_op(χ) — a function of a single state χ relative to the null state 0. This anchored quantity is exactly the plane-capacity C(χ, 0) of the 2-plane (here, the section through χ and 0), so OP10's Vol_op(χ) *is* the sectional capacity, and its identification with $\|\chi\|^2$_ℂ is an area-scaling (dimension-2) claim from the outset — consistent with §3.1's diagonal-anchored remark and with the dimension-2 reading here. Proposition 3.2's "≡" therefore quantifies over the two-state plane-capacity, where the exponent is fixed at 2 by the sectional construction; the L_pack length uses the ambient d_op-capacity; and OP10's object is the former. With the functional pinned, "up to normalization" carries only a constant, not a hidden exponent.

This is the paper's honesty dividend on the inner product. The squaring is **not a new gap introduced here.** It is the gap the programme has carried since PAMV — the packing-volume identification was always conventional — surfacing in sharper form. We do not pretend to close it; we *locate* it as a condition on the reversible symmetry group, note any progress on PAMV Open Problem 10 transfers directly, and record that its linear-and-complex companions, formerly part of the same gap, are now closed.

3.7 Status of the inner-product obstruction

Sub-part of component (i)	Mechanism	Status
Linear carrier (⊕ = amplitude addition)	path-sum construction (Double Square §3.5/§6.1)	**closed (proven)**

Sub-part of component (i)	Mechanism	Status
Complex phase (per-path U(1) phase)	holonomy argument (Physical Necessity §2/App C)	Layer-1 **proven**, Layer-2 **conditional** (§C.8)
Squaring / inner product (conserved norm is ℓ^2)	assumed via "reversible = isometry" (A4)	**open** — relocated to U(d)-invariance ≡ PSD kernel ≡ A7 (≡ PAMV OP 10)

So Obstruction A is **partly closed and sharply located**: two of its three sub-parts are supplied (one proven, one proven-modulo-continuity), and the residue is a single named condition on the symmetry group, equivalent to an inherited open problem. This is strictly more informative than the earlier "assume quadraticity," and it abandons the overclaim — present in the Born-rule papers' abstracts — that "no aspect of quantum mechanics is assumed": the inner product reappears, openly, at the squaring step.

---

4. Sector Decomposition from Symmetry-Compatible Packing

The second ODG component follows from packing plus the inherited substrate symmetry — no obstruction A or B here, but a dependence on H0 that must be charged. Throughout §§4–6 the carrier $\mathcal{A}$_pack is the **complex** carrier supplied by Input H0 (§3.5.2): the seven one-dimensional $\mathbb{Z}_7$ character sectors written below exist as such only over $\mathbb{C}$. This is stated up front because the complexness is *spent* here while being *bought* in §3, and it would be the same burial the paper polices to credit the character decomposition to H1–H2 alone. The real fallback — what survives on H1–H2 *without* H0 — is recorded in Remark 4.3.1.

4.1 Substrate equivariance

Let G be the substrate symmetry group acting on the carrier; in the VERSF setting G = $\mathbb{Z}_7$. Write the action ρ : G → Isom($\mathcal{A}$_pack, ‖·‖_pack), the group of packing-norm isometries — *not* U($\mathcal{A}$_pack), since unitarity presupposes the inner product (the squaring residue), which is not available at the H1–H2 level (Remark 4.1.1). Once the residue resolves, packing-norm isometries coincide with unitaries and ρ lands in U($\mathcal{A}$_pack); until then it lands in the norm-isometry group. The action is **packing-compatible** if

Vol_op(ρ(g)M) = Vol_op(M) for all admissible M and all g ∈ G.

4.2 The packing-equivariance axiom

Axiom FP2 (Packing equivariance). The substrate symmetry action preserves finite packing structure:

$\rho(g) : \Sigma(M) \mapsto \Sigma(\rho(g)M)$, with $|\Sigma(M)| = |\Sigma(\rho(g)M)|$.

Substrate symmetry thus acts by packing-preserving transformations.

4.3 Character-sector decomposition

Because $\mathbb{Z}_7$ is finite abelian, all *complex* irreducibles are one-dimensional characters $\chi_\alpha : \mathbb{Z}_7 \to U(1)$, $\alpha \in \mathbb{Z}_7$ (the one-dimensionality is a fact about $\mathbb{C}$; cf. Remark 4.3.1 for the real case). On the H0 complex carrier the decomposition into character sectors is

$$\mathcal{A}\_pack = \bigoplus_{\alpha \in \mathbb{Z}_7} V_\alpha,\quad V_\alpha = \{ \psi \in \mathcal{A}\_pack : \rho(g)\psi = \chi_\alpha(g)\psi \text{ for all } g \in \mathbb{Z}_7 \}.$$

Proposition 4.1 (Sector decomposition — conditional on H1–H2 for the real decomposition; on H1–H2 + H0 for the seven complex character sectors). Under FP1, FP2, H1–H2, the packing-preserving substrate symmetry action decomposes the carrier into a direct sum of $\mathbb{Z}_7$-irreducible subspaces. Over the H0 complex carrier this is the seven-character decomposition $\mathcal{A}\_pack = \bigoplus_{\alpha\in\mathbb{Z}_7} V_\alpha$. **Status: the real irreducible decomposition is conditional on H1–H2 only; the seven one-dimensional complex character sectors are additionally conditional on H0. Neither depends on Obstruction A or B.**

Proof. FP2 makes $\rho(g)$ packing-volume-preserving. This gives the isometry step directly from the L_pack construction (§3.1), not from Proposition 3.1 (which establishes only that the normed carrier *exists*): since $\rho(g)$ preserves Vol_op on every incremental region $M_{\gamma(t_i),\gamma(t_{i+1})}$ along any admissible path, it preserves each increment $[Vol\_op]^{1/d\_op}$, hence the packing length L_pack of every path, hence the infimum d_pack — so $\rho(g)$ is a d_pack-isometry, i.e. an isometry of $(\mathcal{A}\_pack, \|\cdot\|\_pack)$. Since $\mathbb{Z}_7$ is finite abelian, the representation decomposes into irreducibles. Over the complex H0 carrier the irreducibles are the one-dimensional characters $\chi_\alpha : \mathbb{Z}_7 \to U(1)$, giving the seven eigenspaces V_α; distinct characters are separated by the symmetry action, so the V_α are linearly independent and the decomposition is direct. ■

Remark 4.3.1 (The real fallback, and why the seven-character form needs H0). Without H0 the carrier is, by H2, a real normed vector space (H2 specifies a vector-space structure but not a field, and absent the holonomy argument the natural field is $\mathbb{R}$). Over $\mathbb{R}$ the only one-dimensional irreducible of $\mathbb{Z}_7$ is the trivial character; the six nontrivial complex characters pair as α with $-\alpha$ — that is, {1, 6}, {2, 5}, {3, 4} — into three two-dimensional *real* irreducibles R_k ($k = 1, 2, 3$), the k-th carrying the rotation by $2\pi k/7$. So the real isometric decomposition is

$$\mathcal{A}\_pack^{\mathbb{R}} = V_0 \oplus R_1 \oplus R_2 \oplus R_3,\quad (V_0 \text{ trivial 1-dim, each } R_k \text{ a 2-dim rotation block}),$$

not seven one-dimensional sectors. The seven-character form $\mathcal{A}\_pack = \bigoplus_\alpha V_\alpha$ *is precisely the* ***H0 complexification*** *of this real decomposition: passing to the complex carrier splits each real rotation block R_k into the conjugate character pair $V_\alpha \oplus V_{-\alpha}$*. So the clean abelian character structure that §§4–6 rely on is an H0 result, not an H1–H2 one. We keep the real decomposition as the unconditional (H1–H2-only) content and the seven-character decomposition as its H0 upgrade.

Remark 4.1.1 (Orthogonality is weaker than it looks without the inner product). The decomposition into *irreducible* sectors is *direct* (a vector-space fact, from finite-abelian representation theory) under H1–H2 alone — into the real irreducibles $V\_0 \oplus R\_1 \oplus R\_2 \oplus R\_3$, and into the seven one-dimensional character sectors V_α once H0 supplies the complex carrier (§4.3.1). The direct-sum property holds at both levels; what neither H1–H2 nor H0 supplies is **orthogonality**, which presupposes the inner product, i.e. the squaring residue of Obstruction A. Without it, the sectors are an independent direct-sum decomposition separated by the symmetry; with it, distinct-character subspaces are genuinely orthogonal (an invariant inner product makes inequivalent isotypic components orthogonal). We therefore record sector decomposition as reconstructed at the direct-sum level unconditionally on the squaring (the character form additionally requiring H0), and at the orthogonal-decomposition level conditional on the squaring. For the downstream UMP/Born-rule argument the direct-sum structure plus the inner product (once available) suffices.

---

5. Projection Structure from Maximal Distinguishability Partitions

The third component. Like §4, this goes through under H1–H2, with the orthogonality caveat of Remark 4.1.1 applying again.

5.1 Distinguishability partitions

A finite family $\mathcal{P} = \{M_i\}_{i=1}^n$ is a **distinguishability partition** of M if:

1. the M_i are pairwise operationally disjoint;
2. distinguishability sets in distinct M_i are mutually distinguishable;
3. packing volume is additive: $Vol\_op(M) = \Sigma_i Vol\_op(M_i)$.

5.2 Projection reconstruction

Each distinguishability partition induces a decomposition $\mathcal{A}\_pack = \oplus_i W_i$. The operational projection $P_i : \mathcal{A}\_pack \to W_i$ satisfies

$P_i^2 = P_i, P_i† = P_i, P_iP_j = 0 \ (i \neq j)$,

where the adjoint † and the orthogonality $P_iP_j = 0$ are defined once the inner product is available.

Proposition 5.1 (Projection structure — conditional on H1–H2; orthogonal projections additionally conditional on the squaring residue). If distinguishability partitions are packing-additive and maximal distinguishability sets in distinct partition elements are mutually distinguishable, the carrier admits a partition-induced direct-sum decomposition; given the inner product (squaring residue resolved), this lifts to an orthogonal projection lattice. **Status: direct-sum decomposition conditional on H1–H2; orthogonal-projection lattice additionally conditional on the squaring residue.**

Proof. Packing-additivity (condition 3) implies the distinguishability content in M_i and M_j does not overlap. To pass from non-overlap to *linear independence* of the subspaces W_i — not merely their set-theoretic disjointness — we use Input H0: under the path-sum picture each W_i is spanned by the amplitude contributions of the path-set resolving M_i, and disjoint path-sets contribute amplitudes along independent directions, so packing-disjoint partition elements give linearly independent W_i and the direct sum $\bigoplus_i W_i$. (Absent H0, non-overlap of distinguishability content does not by itself force linear independence of the carrier subspaces.) Non-overlap corresponds to orthogonality *once an inner product is present*: with it, the W_i are mutually orthogonal closed subspaces and the orthogonal-projection theorem supplies a unique self-adjoint idempotent P_i onto each. Absent the inner product, the construction yields the direct-sum decomposition and idempotents P_i with $P_i^2 = P_i$, but self-adjointness and $P_iP_j = 0$ are not defined. ■

This recovers ODG component (iii), with the same dependence of "orthogonal" on the squaring that §4 carries. Component (iii) does **not** depend on Obstruction B.

---

6. Admissible Transport from Packing-Preserving Automorphisms

The fourth component. Clean, modulo the now-familiar orthogonality caveat.

6.1 Packing-preserving, substrate-compatible automorphisms

An automorphism U of the carrier is **packing-preserving** if Vol_op(UM) = Vol_op(M) for every admissible M, and **substrate-compatible** if $U\rho(g) = \rho(g)U$ for all $g \in \mathbb{Z}_7$.

6.2 The transport group

Define

$U\_adm(\mathcal{A}\_pack) = \{ U \in GL(\mathcal{A}\_pack) : U$ preserves packing and commutes with $\rho(\mathbb{Z}_7) \}$.

Commutation with the substrate action forces U to preserve character sectors, $U(V\_\alpha) \subseteq V\_\alpha$, whence

$U_{adm}(\mathcal{A}_{pack}) = \prod_\alpha GL(V_\alpha)$ (and $= \prod_\alpha U(V_\alpha)$ once the inner product makes packing-preservation = unitarity).

Proposition 6.1 (Admissible transport — conditional on H1–H2 + H0; the unitary form additionally conditional on the squaring residue). Packing-preserving, substrate-compatible automorphisms reconstruct the block-diagonal transport group $\prod_\alpha GL(V_\alpha)$; given the inner product, packing-preservation coincides with unitarity and the group is $\prod_\alpha U(V_\alpha)$. **Status: the block-diagonal form over the seven complex character sectors is conditional on H1–H2 + H0 (the indexing is over the H0 character decomposition of §4.3; on the H1–H2-only real decomposition the analogous statement is block-diagonal over $V_0 \oplus R_1 \oplus R_2 \oplus R_3$ with GL/rotation-commutant blocks, Remark 4.3.1); the unitary form is additionally conditional on the squaring residue.**

Proof. Substrate compatibility forces commutation with the $\mathbb{Z}_7$ action; by Schur over the character sectors, any such automorphism is block-diagonal with respect to $\bigoplus_\alpha V_\alpha$, giving $\prod_\alpha GL(V_\alpha)$. With the inner product, packing-preservation = norm-preservation = unitarity, so each block lies in $U(V_\alpha)$. ■

This recovers ODG component (iv). Note the circularity-free reading enabled by §3.5.3: the transport group's *unitarity* and the inner product's *existence* are the same residue (face 1) — admissible transport being unitary just *is* the statement that the reversible normalization-preserving group is U(d). Independent of Obstruction B; dependent on the squaring residue only for the upgrade from GL to U.

---

7. Refinement Structure from Resolution Stability — and the Refinement Obstruction

The fifth component is the one the ODG paper's own sub-link analysis (its Open Problem 1, sub-link ii.f) flagged as the structural risk. It is Obstruction B, and unlike Obstruction A it is genuinely new to this reduction.

7.1 Refinement of distinguishability partitions

Let $\mathcal{P} = \{M_i\}$ be an admissible distinguishability partition of M. A **refinement** $\mathcal{P}' = \{M_{ij}\}$ satisfies $M_i = \bigcup_j M_{ij}$ and packing-additivity $Vol_{op}(M_i) = \Sigma_j Vol_{op}(M_{ij})$. A refinement is **resolution-stable** if it remains admissible under rescaling of operational resolution $\Delta_{op} \mapsto \lambda\Delta_{op}$ for all allowed λ.

7.2 The refinement-stability hypothesis — stated as an obstruction

Hypothesis FP3 (Refinement = resolution-stable packing). An orthogonal decomposition of an admissible subspace is physically admissible *if and only if* it arises from a resolution-stable refinement of a packing-additive distinguishability partition.

We deliberately do **not** label FP3 an axiom on a par with FP1–FP2. FP1–FP2 are about *packing* (counting and symmetry); FP3 is a bridge identifying an *admissibility class* (which decompositions count, in the PAMV §6.2 sense) with a *packing class* (resolution-stable refinements). That bridge is exactly the content the ODG paper inherited from PAMV §6.2 as a structural input. FP3 proposes to ground it in packing — but proposing is not proving.

Obstruction B (Refinement gap). FP3 does not follow from FP1 and FP2. The forward direction (every resolution-stable packing refinement is admissible) is plausible from packing-additivity. The reverse direction (every admissible decomposition is resolution-stable packing refinement) is the load: it asserts that PAMV §6.2's admissibility apparatus introduces *no* admissible decompositions beyond those visible to packing-and-resolution. Nothing in FP1–FP2 rules out an admissible decomposition that packing-and-resolution cannot see.

Proposition 7.1 (Refinement structure — conditional on FP3). Under FP3 (and H1–H2), refinement structure is reconstructed as the class of resolution-stable packing decompositions together with the decomposition-admissible unitaries W relating them. **Status: conditional on FP3 ≡ Obstruction B.**

Proof. Under FP3, admissible decompositions are exactly the resolution-stable packing refinements. These are packing-additive, so Proposition 5.1 assigns each an (orthogonal, given the squaring residue) decomposition. A unitary $W : V_\alpha \oplus V_\beta \to V_{\alpha'} \oplus V_{\beta'}$ between two such decompositions preserves the refinement relation iff it maps resolution-stable partition elements to resolution-stable partition elements; these are exactly the decomposition-admissible (U-dec) transformations. ■

7.3 Why Obstruction B is independent of Obstruction A

The two obstructions are logically independent, and it is worth saying why, because it determines the branch structure of §8.

Obstruction A's residue is about the *quadraticity of the normalization* (is the conserved reversible-invariant norm the squared Hilbert norm?). Obstruction B is about the *completeness of the admissibility class* (does packing-and-resolution see every admissible decomposition?). One can imagine resolving either without the other: a Hilbertian carrier whose admissible decompositions exceed the resolution-stable ones (A resolved, B open), or a non-Hilbertian normed carrier on which admissible = resolution-stable refinement holds exactly (B resolved, A open). Neither obstruction's resolution supplies the other's. This independence makes the §8 assembly a genuine 2×2 of branches rather than a single chain.

Remark 7.3.1 (Relation to the ODG paper's sub-link ii.f). Obstruction B is precisely the difficulty the ODG paper isolated as sub-link (ii.f) of its Open Problem 1 — the refinement component's resistance to reduction because the (U-dec) class and the sector-attachment convention carry structural commitments not visibly reducible to Vol_op plus $\mathbb{Z}_7$-equivariance. The present paper does not resolve ii.f; it gives it a sharper name (FP3), an explicit if-and-only-if statement, and an account of which direction (the reverse) carries the load.

---

8. Assembly: The ODG Reduction Theorem

We now assemble, refusing the temptation to call the result a derivation.

8.1 The component ledger

ODG component	Reconstructed from	Conditional on
(i) carrier + inner product	linear carrier + complex phase: path-sums over U(1) holonomy (§3.5.2), i.e. **Input H0**; inner product: the squaring residue	**linear carrier proven given H0**; **phase** conditional on H0 + holonomy continuity upgrade; **inner product** on the squaring residue (≡ PAMV OP 10)
(ii) sector decomposition	packing-equivariant $\mathbb{Z}_7$ symmetry (§4)	H1–H2 for the real irreducible decomposition; **+ H0** for the seven one-dimensional complex character sectors (over $\mathbb{R}$: trivial $\oplus$ 3·2-dim, Rmk 4.3.1); orthogonality: + squaring residue
(iii) projection structure	packing-additive partitions (§5)	H1–H2; **linear independence of the W_i uses H0** (see Prop 5.1); orthogonality: + squaring residue
(iv) admissible transport	packing-preserving substrate automorphisms (§6)	H1–H2 **+ H0** (block structure indexed over the complex character sectors); unitarity: + squaring residue
(v) refinement structure	resolution-stable packing refinements (§7)	**Obstruction B** (≡ FP3)
(vi) finite packing	primitive input (FP1)	— (axiom)

Reading the ledger: (vi) is free. Components (ii)–(iv) reconstruct under the mild H1–H2 *as real irreducible decompositions*, with their metric content unconditional on the obstructions; their **seven-character / complex form** additionally requires H0, and their *orthogonal/unitary* content waits on the squaring residue. Component (i) is split: its linear-and-complex carrier is supplied by **Input H0** (the phase modulo the continuity upgrade), its inner product waits on the squaring residue. Component (v) waits on Obstruction B alone. The lone non-packing primitive in the ledger is H0, and it is worth noting that it enters for three *distinct* reasons rather than as a catch-all: on row (i) as the **carrier** itself (linearity + complex phase), on rows (ii)/(iv) as **complexness** (the seven-character / unitary-block form), and on row (iii) as **subspace-independence** (linear independence of the partition subspaces W_i, Prop 5.1). Each is a genuine, separately-identifiable use; H0 is named on every row that uses the complex carrier — (i), (ii), (iii), (iv) — precisely so that neither the carrier's linearity *nor the abelian character structure §§4–6 run on* is silently credited to finite packing.

8.2 The theorem, stated as a conditional

Theorem 8.1 (ODG Reduction Theorem). Assume FP1, FP2, H1, H2, and Input H0 (the path-sum substrate, §3.5.2). Then:

1. **(Unconditional on the two obstructions.)** Finite packing, *together with H0*, reconstructs the direct-sum carrier — its *linear* structure and *complex phase* supplied constructively by the path-sum / U(1)-holonomy route (the phase modulo the Layer-2 continuity upgrade) — together with the sector decomposition (ii at direct-sum level), partition-induced subspace structure (iii at direct-sum level), block-diagonal transport ∏_α GL(V_α) (iv at GL level), and packing measure (vi). The linearity of this part rests on H0, not on FP1–FP2; "unconditional on the obstructions" does not mean "from finite packing alone."
2. **(Conditional on the squaring residue ≡ PAMV Open Problem 10.)** Additionally, the carrier carries an inner product (i complete); the sector decomposition, projections, and transport upgrade to their orthogonal/unitary forms.
3. **(Conditional on Obstruction B ≡ FP3.)** Additionally, the refinement structure (v) is reconstructed.
4. **(Both obstructions resolved.)** All six components of ODG are reconstructed, and the carrier carries a full ODG in the sense of the ODG paper.

Status: Part 1 is **conditional on H0 and H1–H2** (proven from FP1–FP2, the inherited completion hypotheses, and the path-sum substrate H0 of the Born-rule papers), with the complex phase additionally conditional on the holonomy continuity upgrade. Parts 2 and 3 are **conditional** on their named obstructions, which are **open**. Part 4 is the conjunction. At no point is the full six-component reconstruction asserted as **proven**, and at no point is the linear carrier credited to finite packing without H0.

Proof. Part 1 collects the unconditional halves of Propositions 4.1, 5.1, 6.1, the axiom FP1, and the constructive carrier of §3.5.2 (which assumes H0; Propositions 4.1 and 5.1 inherit the H0-dependence noted in their statements). Part 2 adds Proposition 3.2 (the squaring residue supplies the inner product) and the orthogonality/unitarity upgrades of Remark 4.1.1 and Propositions 5.1, 6.1. Part 3 adds Proposition 7.1. Part 4 is Parts 2 and 3 together. The independence of the obstructions (§7.3) ensures Parts 2 and 3 are separately attainable. ■

8.3 What the theorem is and is not

It **is** a map of the reduction: four components down to packing, a constructively-supplied linear-and-complex carrier, one residue (the squaring) blocked by an inherited problem, one component (v) blocked by a new hypothesis. It **is not** a derivation of ODG from packing — that statement is true only in Part 4, conditional on two open problems, one of which (the squaring) the programme has carried since PAMV. The honest one-line summary:

Finite packing, *together with the path-sum substrate H0* supplying the Born-rule papers' constructive carrier, reconstructs ODG outright except for the squaring of the inner product and the completeness of the refinement class — two independent, individually-isolable obstructions.

---

9. Consequence for the Born Rule

The ODG paper established

ODG → OIP → ODG-compatibility → UMP → Born rule.

The present paper supplies the candidate preceding link, conditionally:

finite packing (+ H0) →(Parts 1–4 of Theorem 8.1) ODG.

Corollary 9.1 (Conditional finite-packing route to the Born rule). Given the ODG-to-Born reconstruction of the prior paper, and Theorem 8.1 Part 4 (both obstructions resolved), the Born rule is traceable to finite distinguishability packing together with the path-sum substrate:

finite packing (+ H0) → ODG → OIP → ODG-compatibility → UMP → Born rule.

Status: conditional on Obstruction A's squaring residue (≡ PAMV Open Problem 10) and Obstruction B (≡ FP3), both open.

Proof. Under both obstructions, Theorem 8.1 Part 4 reconstructs ODG; the ODG paper then yields ODG-compatibility ⟹ UMP ⟹ Born rule. ■

Remark 9.1.1 (The honest weakening, and the asymmetry between the conditions). Corollary 9.1 is multiply conditional and we do not overstate it. Note first the asymmetry on the inner-product side. The squaring residue is *not specific to this paper* — it is PAMV Open Problem 10, which the Born-rule programme has depended on since the packing-volume identification was made conventional. So Corollary 9.1 inherits that dependency rather than adding it, and the *companions* of the residue that earlier framings lumped with it — the carrier's linearity and complex phase — are now closed (§3.5.2), so the inner-product debt is genuinely smaller than before: exactly the squaring, exactly the U(d)-invariance / PSD-kernel / A7 condition, nothing more. But honesty requires naming a second debt that does *not* reduce to a prior open problem: **Input H0, the path-sum substrate.** The constructive closure of the linear-and-complex carrier rests on H0, which is independent of FP1–FP2 and not derived here. So the fair statement of marginal contribution is not "the finite-packing route is complete except for FP3," but rather: *given the path-sum substrate H0, and modulo the squaring the programme already owed (A), the route to the Born rule is complete except for the single new hypothesis*

FP3 (B). There are thus two debts beside the inherited squaring — the new FP3 (B), and the imported H0 — and the marginal-contribution claim is earned only once both are on the table. With them named, the claim is sharper and far more defensible than "finite packing yields the Born rule"; without H0 named, it would be the very burial the paper sets out to refuse.

---

10. The Two Obstructions, Compared

Because the obstructions are the paper's real content, we compare them directly.

	Obstruction A (the squaring residue)	**Obstruction B (refinement)**
What it asserts	the conserved reversible-invariant normalization is the squared Hilbert norm — equivalently, the reversible normalization-preserving group is U(d) (≡ PSD kernel ≡ A7 bilinearity)	admissible decomposition = resolution-stable packing refinement
Already partly closed?	**Yes** — the linear carrier (proven) and complex phase (proven modulo continuity) are supplied (§3.5.2); only the squaring remains	no sub-parts closed
ODG component blocked	the *inner product* of (i), plus orthogonal/unitary upgrades of (ii)–(iv)	(v)
New to this paper?	**No** — the residue ≡ PAMV Open Problem 10	**Yes** — sharpens ODG-paper sub-link (ii.f)
Where the residue sits	the symmetry group on outcomes (U(d) vs other isometry group); negative type of D is its downstream shadow on the relational metric	reverse inclusion (no admissible decomposition invisible to packing)
Plausible attack	exhibit a substrate-admissible reversible dynamics whose conserved outcome normalization is *not* the ℓ^2 norm, or prove the substrate forces "reversible = isometry of a $\mathbb{C}$ inner product" (§11.1)	derive FP3-reverse from packing-additivity + resolution-stability (§11.2)
Cost if it fails	if the squaring is merely unprovable: carrier stays linear-and-complex but normed-not-Hilbert, and the Hilbert structure stays conventional. If a non-ℓ^2 conserved normalization is actually *constructed* (§11.1's counterexample branch): the squared-norm identification is **proven false**, PAMV OP 10 settles negatively, and the inner product is established as independent input — not "left conventional" but positively excluded	refinement stays an independent admissibility input; (i)–(iv),(vi) still reduce

Two structural observations:

First, **Obstruction A's residue is the deeper of the two**, because it is upstream of orthogonality, self-adjoint projections, and unitarity — the entire Hilbert character of the carrier. But it is the one this paper is *least* on the hook for: it predates this paper (PAMV Open Problem 10), and its linear-and-complex companions are now closed, so what remains is a single sharp condition on the reversible symmetry group, and any solution at PAMV OP 10 flows here automatically (Proposition 3.2).

Second, **Obstruction B is the one this paper genuinely owes**, and it is more local: its failure costs only component (v), leaving a clean five-of-six reduction (with the squaring caveat) intact. The *new* frontier this paper opens is exactly FP3 — can the reverse inclusion be derived? — while Obstruction A is a frontier the programme already knew it had to cross.

---

11. Limitations and Open Problems

11.1 Is the squaring forced? — the residue of Obstruction A

The sharpest form of Obstruction A, after §3.5: the linear carrier and complex phase are supplied; the open question is whether the conserved, reversible-transport-invariant normalization on $\mathbb{C}^d$ is the ℓ^2 norm — equivalently, whether the reversible normalization-preserving symmetry group is U(d) preserving a positive-definite complex inner product (≡ PSD of the correlation kernel ≡ bilinearity of selection, §3.5.3). The Born-rule papers obtain this only by *assuming* "reversible = isometry of the distinguishability metric" and then demanding invariance of the normalization under the resulting group; they do not derive that the group is unitary on $\mathbb{C}^d$ rather than the isometry group of some other outcome metric.

The decisive test, at the group level. The per-path holonomy argument delivers U(1) phase per path; the residue is the jump to U(d) on the outcome vector. The test is therefore the group-level analogue of the old ℓ^p probe:

- **Attempt the counterexample.** Construct a substrate-admissible reversible dynamics whose conserved outcome-vector normalization is *not* the ℓ^2 norm — e.g. an isometry group of a non-Euclidean outcome metric (an ℓ^p, $p \neq 2$, outcome norm) consistent with TPB, BCB, temporal extensibility, and the path-sum carrier. If such a dynamics exists, the squaring is genuine extra structure: Obstruction A's residue is *irreducible*, PAMV Open Problem 10 resolves negatively (the packing-volume / squared-norm identity is stipulation, not theorem), and the inner product is proven independent input. This is a sharp, citable result, worth more than any closure-by-renaming.
- **If the counterexample provably cannot be built.** Then the obstruction to its construction — the specific way TPB / BCB / temporal extensibility / the path-sum carrier force the reversible group to be U(d) — *is* the derivation of the inner product, hence of the squared-amplitude Born rule, from inputs that are not quadraticity assumed.

Either outcome settles the question; neither is a reformulation. The relational sub-question of §3.5.1 (do (D1)–(D5) force Schoenberg negative type?) is the shadow of this test one level down, on the metric rather than the group; a negative-type counterexample (an ℓ^p relation satisfying (D1)–(D5)) and a non-ℓ^2-normalization reversible dynamics are two faces of the same construction. We record our weak prior unchanged: we expect symmetry + identity + triangle-type composition to be *insufficient*, with the weight (if carriable at all) on the conservation/recombination structure — which may itself be equivalent to the U(d) condition, in which case the honest conclusion is irreducibility with a sharply named residue. This is the highest-leverage open problem touching the paper, shared with PAMV Open Problem 10.

11.2 Can FP3 be derived from FP1–FP2? (Obstruction B)

The main *new* open problem. The forward direction of FP3 (resolution-stable packing refinements are admissible) is plausible; the reverse (admissible decompositions are exhausted by resolution-stable packing refinements) is the question. A derivation would need to show that PAMV §6.2's admissibility apparatus introduces no decompositions invisible to packing-and-resolution — equivalently, that the (U-dec) class is generated by resolution-stable refinements. If instead a counterexample exists (an admissible decomposition not resolution-stable), FP3 fails and refinement remains an independent input, consistent with the ODG paper's flagging of sub-link (ii.f) as the structural risk.

11.3 The phase-side continuity upgrade

Distinct from the squaring residue, and on the carrier side rather than the inner-product side: promoting the holonomy argument's Layer 1 (finite holonomy impossible — proven) to Layer 2 (continuous U(1) holonomy) is the programme's own four-item obligation (Physical Necessity §C.8): operational-topology formalization; the nearly-canceling-paths accumulation argument; dimensionality $d_H = 1$ from TPB irreducibility; compactness of the holonomy connected component. Until discharged, the *complex phase* of the carrier is conditional (though the *linearity* is not). This is conjectural-to-conditional with a named, structured roadmap.

11.4 Independence of the obstructions

We argued (§7.3) that A and B are logically independent. A sharper result would *prove* independence by exhibiting carriers realizing each of the three non-trivial branches (A-only, B-only, neither). Constructing an A-resolved/B-open carrier seems most tractable (take a genuine Hilbert carrier and enlarge its admissible-decomposition class by fiat); the B-resolved/A-open carrier is harder and would itself be a useful object.

11.5 Exhaustiveness of the six components

The reduction inherits from the ODG paper the convention that the six components exhaust operational distinguishability geometry. If a seventh operational structure exists, the reduction is incomplete in a way orthogonal to A and B. We do not revisit this convention here; it is flagged as inherited.

11.6 Empirical handle

We are deliberately conservative here. If finite packing genuinely co-determines ODG, then probability structure and operational packing volume are structurally linked rather than correlated by construction, and one expects *some* empirical signature of that link — a deviation from packing-volume invariance ought to induce a corresponding deviation in admissible probability structure. But we do not currently have a sharp empirical handle: we cannot yet name the observable, the regime, or a null hypothesis precise enough to constitute a test, and the ODG paper's "ODG-test 2," to which such a handle would connect, is not developed to that level either. The honest statement is therefore that the structural link *implies* a measurable consequence should exist while leaving its form open; converting that implication into a concrete observable — most naturally on the resolution of the squaring residue, where the probability/volume identification stops being conventional — is itself an open problem, not a result claimed here. We flag it as the empirical counterpart to §11.1 rather than as a prediction.

11.7 Do H0 and finite packing share a substrate root?

This paper treats finite packing (FP1–FP2) and the path-sum substrate (H0) as two independent structural primitives (§3.5.2). That independence is what makes the reduction honest, but it also opens the obvious next question: are the two primitives genuinely independent, or do both descend from a single deeper structure — in which case the reduction would have *one* substrate input rather than two, and the title's "from finite packing" would be closer to literal than to shorthand?

There is a natural candidate common root inside the programme. Finite packing is plausibly a statement about the **tick-per-bit (TPB) and bit-conservation (BCB)** structure — bounded distinguishability capacity per region is what a finite commitment-rate substrate would produce. The path-sum substrate H0 is *also* sourced, in the Born-rule papers, from TPB plus the commitment/holonomy structure (the U(1) phase comes from TPB time, §3.5.2 Layer 1). So both primitives gesture at the same TPB/BCB substrate. If that common sourcing can be made into a derivation — *finite packing and the reversible path-sum substrate are two faces of one TPB/BCB commitment structure* — then H0 is not an independent import but a co-consequence, and the two debts of Remark 9.1.1 collapse toward one.

We do **not** establish this here, and we flag two reasons for caution rather than asserting the collapse. First, the two primitives constrain different things — packing bounds *capacity*, H0 supplies *reversible composition and phase* — and a common informal source is not a derivation that one structure yields both. Second, the holonomy continuity upgrade (§11.3) is itself still open on the H0 side, so even the TPB → H0 half is not fully discharged. The precise open problem is therefore: *exhibit a single substrate-level structure (presumably the TPB/BCB commitment bath) from which both FP1–FP2 and H0 provably follow, or show they are independent by constructing a model satisfying one and violating the other.* A positive result would reduce the paper's two substrate inputs to one; a negative result would establish that capacity-boundedness and reversible-phase composition are genuinely separate facts about the substrate. This is the most natural reduction target the present accounting opens, and it is distinct from both the squaring residue (§11.1) and the phase-side continuity upgrade (§11.3).

---

12. Conclusion

The ODG paper showed that the Born rule is the unique invariant measure compatible with Operational Distinguishability Geometry, with ODG itself a structural input. This paper asks whether ODG descends from finite distinguishability packing, and answers: **partly, and the parts that resist are worth more than the parts that go through.**

What goes through: from finite packing plus substrate equivariance, four of the six ODG components — sector decomposition, projection structure, admissible transport, and packing itself — are reconstructed under mild completion hypotheses (§§4–6). Packing volume induces a metric; the metric induces a norm; equivariant symmetry induces sectors; additive partitions induce projections; packing-preserving automorphisms induce transport. And the carrier itself — formerly part of the inner-product obstruction — is now supplied *constructively*: its linear structure as path-sums (proven) and its complex phase as U(1) holonomy (proven modulo a flagged continuity upgrade), drawing on the programme's Born-rule papers.

What resists, and where the paper's content lives:

- **The squaring of the inner product (Obstruction A, residue).** A counting-measure norm is not automatically quadratic, and — the structural lesson of Remark 3.4.0 — the *unary diagonal* packing norm cannot generate the *binary* parallelogram law without re-importing quadraticity (the binary content lives instead in Vol_op over two-state regions, which is where §3.5 poses the question). The Born-rule papers' constructive carrier closes the linear and complex parts of the inner-product gap; what remains is the squaring, which every route reaches only by assuming the reversible normalization-preserving group is the unitary group U(d) preserving a positive-definite complex inner product (≡ PSD of the correlation kernel ≡ bilinear selection). This is *assumed* via "reversible = isometry," not derived. We show it *is* PAMV Open Problem 10 (Proposition 3.2) — inherited, not new — with the classical negative-type / Schoenberg condition as its downstream shadow on the relational metric (§3.5.1). We decline to claim closure of the squaring, and we decline the Born-rule papers' abstract claim that "no aspect of quantum mechanics is assumed": the inner product reappears, openly, at this step.
- **The refinement structure (Obstruction B).** Recovering refinement requires identifying admissible decomposition with resolution-stable packing refinement (FP3). We show FP3 does not follow from the packing axioms, isolate the reverse-inclusion direction as the load, and connect it to the ODG paper's sub-link (ii.f). This is the genuinely new frontier the paper opens.

The result is therefore a **reduction with a constructive carrier and two isolated, independent obstructions**, not a derivation:

Finite packing reconstructs ODG components (ii)–(iv) at the real-decomposition level and (vi); with the path-sum substrate H0 it additionally supplies the linear-and-complex carrier of (i) and the seven-character / complex form of (ii)–(iv); the inner product of (i) follows iff Obstruction A's squaring residue (≡ PAMV OP 10) resolves; component (v) iff Obstruction B (≡ FP3) resolves.

Conditional on both, the VERSF hierarchy deepens by one layer:

finite packing (+ H0) → ODG → OIP → ODG-compatibility → UMP → Born rule,

tracing the Born rule, conditionally, to the finite packing structure of distinguishable states *together with the path-sum substrate*. Conditional on Obstruction B alone, the route is complete modulo the squaring the programme already owed since PAMV. Conditional on neither, a four-of-six reduction, a constructive carrier, and a precise localization of the residue still stand.

The principal unresolved questions are sharp and independently attackable: *is the conserved reversible-invariant normalization the ℓ^2 norm — equivalently, is the reversible symmetry group U(d) — or is this irreducible structure assumed via "reversible = isometry"* (Obstruction A's squaring residue / PAMV Open Problem 10), and *can the reverse inclusion of FP3 be derived from finite packing* (Obstruction B). For the squaring the recommended first move is a construction problem (§11.1): build a substrate-admissible reversible dynamics whose conserved outcome normalization is *not* the ℓ^2 norm. If it exists, the inner product is proven independent structure (a negative settlement of PAMV Open Problem 10, itself a result); if it provably cannot be built, the obstruction to building it *is* the derivation of the squared-amplitude Born rule. Closing Obstruction B would make refinement structure a consequence of bounded operational resolution. Either settlement moves the programme measurably closer to a first-principles account of quantum probability; neither is claimed here.