Operationally Invisible Sector Boundaries

Can Finite Distinguishability Permit a Sector Separation of Admissible Motion?

Keith Taylor VERSF Theoretical Physics Programme versf-eos.com

---

General Reader Summary

The previous papers reduced the origin of quantum probability to a surprisingly specific question.

Reality appears to admit two kinds of reversible motion before a fact is formed.

The first is phase motion: probabilities rotate and interfere while the underlying carving of possibilities stays fixed.

The second is refinement motion: the carving itself changes.

The previous paper established that these two motions belong to the same admissibility class, and that the only remaining obstacle to deriving the Born rule is whether they inhabit one connected space of motion or two separate, walled-off sectors.

This paper asks a sharper version of that question.

A wall between two sectors is a real feature of reality only if something could, in principle, detect which side of the wall a motion is on. The substrate has a smallest scale at which it can tell anything apart at all — a distinguishability floor. So the question becomes: if the substrate cannot tell a phase motion from a refinement motion using any operation available to it, on what grounds could reality place them in different sectors?

We argue that the substrate's own operational standard for what counts as real already forbids walls that nothing can detect. The remaining work then collapses to a single, well-posed problem: showing that phase and refinement motion are operationally indistinguishable not only step by step, but also for finite paths built by composing many steps. We make precise exactly what must still be proven, and what now follows for free.

The same standard turns out to govern probability itself. If the substrate cannot tell two situations apart, then no honest probability can treat them as different. We show in the later sections that this single idea — reality contains only what the substrate could in principle detect — is the common root of two results: that a witnessless wall between sectors is not real, and that probability cannot distinguish what the substrate cannot distinguish. Two principles that previous work treated as separate postulates turn out to be the same principle pointed at two different objects.

Contents

Abstract

Previous work reduced the Born-rule residue to Reversible Connectedness (RC), then decomposed RC into a static component (connectedness of the admissible refinement space) and a dynamical component (Traversability). The static component was reduced to capacity geometry, leaving one unresolved issue: whether admissible reversible motion possesses a genuine sector separation between phase and refinement directions.

This paper formalizes that question and reorganizes the residue.

A topology is introduced on the refinement space $\mathscr{R}$ via a transport metric on capacity partitions, and the admissible-motion space $\mathcal{A}$ is defined together with a projection $\pi : \mathcal{A} \to \mathscr{R}$. The static question becomes whether continuous paths in $\mathscr{R}$ admit admissible reversible lifts in $\mathcal{A}$ (the Admissible Lift Property, ALP).

We make three changes to the prior framing.

First, we reclassify the Operational Sector Principle (OSP). Rather than an independent conjecture, OSP is shown to be a consequence of the VERSF operational ontology already assumed throughout the programme: an admissibility distinction that admits no operational witness is not a structural feature of the substrate. OSP therefore carries the epistemic status [Conditional-on-OO], not [Conjectural].

Second, we sharpen the Fiber–Base Identification (FBI). We separate a sub-floor component — phase and refinement displacements below the distinguishability floor produce no distinguishable substrate states — which is [Conditional/near-trivial], from a composite component — no protocol built from floor-resolved operations recovers the sector distinction for finite paths — which is the genuine [Conjectural] residue, denoted FBI-comp.

Third, we close a gap in the previous chain. Connectedness of the total space does not by itself imply path-lifting. We isolate the missing premise as a Uniqueness-of-Obstruction lemma (UO) — the sector boundary is the sole obstruction to ALP — and carry it explicitly as [Conditional] rather than leaving it implicit.

The net result is that the Born-rule residue reduces, under the operational ontology, to two clearly labelled items: the composite identification FBI-comp and the obstruction-uniqueness premise UO, together with the previously isolated phase-continuity tier. The principal open problem is the derivation of FBI-comp from finite distinguishability.

Finally, we exhibit the architectural payoff. The operational ontology (OO) has a second corollary alongside OSP. Finite distinguishability — via finite packing — generates an operational quotient of the state carrier through a resolution map onto a finite operational state set; operational distinguishability geometry (ODG) is the geometry of that quotient; and the Operational Indistinguishability Principle (OIP) — that any operationally meaningful probability assignment is constant on operationally identical states — is then a theorem about quotient spaces rather than a postulate. OSP and OIP are shown to be the same principle (OO) applied to two objects: distinctions (sector boundaries) and functions (probability assignments), though with different epistemic standing. This embeds the probability programme and the sector programme in one foundation and recovers a strengthened conceptual hierarchy. The two threads are not rival derivations of the Born rule: RC is the connectivity precondition that the measure-uniqueness (UMP) step of the probability chain consumes, as made precise in §12.1.

---

1. Introduction

The previous paper reduced RC to a lifting problem. Let:

$\mathscr{R}$ = admissible refinement space

$\mathcal{A}$ = admissible reversible pre-commitment motion space

$\pi : \mathcal{A} \to \mathscr{R}$

The remaining question was stated as:

Does every continuous path in $\mathscr{R}$ admit an admissible reversible lift in $\mathcal{A}$?

This paper does three things. It frames the lifting question in terms that respect the finite distinguishability floor rather than borrowing continuum infinitesimals. It identifies which of the previously asserted implications are in fact load-bearing and supplies the missing premises explicitly. And it reclassifies the principles in play according to their true epistemic status, so that the residue is neither overstated nor understated.

We adopt the programme's standing labelling convention. Each substantive claim is tagged [Proven], [Conditional], [Conditional-on-OO], or [Conjectural], where [Conditional-on-OO] marks a claim that follows from the VERSF operational ontology stated in §4.

---

2. Refinement Space and Motion Space

We recall the substrate primitives needed here. A refinement R is a capacity partition: an admissible carving of the substrate's distinguishable possibilities at finite capacity. Capacity displacement measures the amount of reallocation across such a partition, in the units fixed by the distinguishability floor.

Definition (transport metric). For refinements R and R′,

d(R, R′) = inf over admissible transport plans T : R → R′ of the total capacity displacement of T.

Three properties require comment, since the topology of $\mathscr{R}$ rests on them. Two of them are genuine closure assumptions on admissibility, stated as such.

(i) **Degeneracy structure.** d(R, R′) = 0 iff R and R′ are floor-indistinguishable. The floor, not exact equality, is the zero of the form. Consequently d is in general a **pseudometric**, not a metric: distinct raw refinements separated only below the floor sit at distance zero. We take $\mathscr{R}$ to be the **raw** refinement space — points are carvings, not floor-classes of carvings — so that $\mathscr{R}$

carries a pseudometric and a non-Hausdorff topology. The operational refinement space is the floor-quotient $\mathscr{R}/\approx\_\mathscr{R}$ obtained by collapsing the zero set of d; that quotient is performed once, in §10, and §2 does not pre-quotient. (This bookkeeping prevents the double-count flagged in §10.2: the Kolmogorov quotient of §10 *is* the floor identification, not a second one.)

(ii) **Symmetry [closure assumption: admissibility is closed under reversal].** If admissible transport is closed under reversal in the pre-commitment regime, the minimal-displacement plan from R to R′ has equal cost to its reverse, giving $d(R, R') = d(R', R)$. Reversal-closure is the defining feature of the pre-commitment regime, but it is an assumption about admissibility, not a property of capacity displacement alone, and we label it as such.

(iii) **Triangle inequality [closure assumption: admissibility is closed under concatenation].** If the concatenation of two admissible transport plans is itself admissible, and capacity displacement is sub-additive under concatenation, then $d(R, R'') \leq d(R, R') + d(R', R'')$. Concatenation-closure is again a closure assumption on admissibility, reasonable in the pre-commitment regime but not automatic, and labelled accordingly.

One disambiguation is essential. The infimum is taken over **admissible** transport plans only. Admissibility is therefore already present in the base topology. This is intended, but it has a consequence we flag now and use in §7: a path that is continuous in $\mathscr{R}$ is a path each of whose steps is realizable by admissible transport in the base. The lifting question is then not whether motion in $\mathscr{R}$ is admissible — that is built in — but whether a base-admissible path can be **carried reversibly** in the full motion space $\mathcal{A}$ without leaving the admissible reversible regime.

Definition (motion space). $\mathcal{A}$ is the space of admissible reversible substrate configurations, equipped with the projection

$$\pi : \mathcal{A} \to \mathscr{R}$$

sending each configuration to the refinement it instantiates. The fiber over R, written $\pi^{-1}(R)$, is the space of phase configurations compatible with the carving R.

---

3. The Admissible Lift Property

Definition (admissible lift). A continuous path γ in $\mathscr{R}$ possesses an admissible lift if there exists a continuous admissible reversible trajectory $\tilde{\gamma}$ in $\mathcal{A}$ with

$$\pi \circ \tilde{\gamma} = \gamma.$$

Definition (ALP). Every continuous path in $\mathscr{R}$ admits an admissible reversible lift.

The previous paper stated an equivalence "RC $\Leftrightarrow$ ALP once $\mathscr{R}$-connectedness has been established." That statement is doubly defective: path-connectedness of the total space $\mathcal{A}$

requires connectedness of the base, lifting of base paths, **and** connectedness within fibers; and the converse direction does not hold without already assuming what is at issue. We therefore state only the implication the downstream chain uses.

Proposition 3.1 [Proven].

($\mathscr{R}$-connectedness ∧ ALP ∧ fiber connectedness) ⟹ RC.

Proof. Given two configurations a, b ∈ $\mathcal{A}$, project to R_a, R_b ∈ $\mathscr{R}$. $\mathscr{R}$-connectedness gives a continuous path γ from R_a to R_b; ALP lifts it to $\tilde{\gamma}$ from some a′ ∈ π^{-1}(R_a) to some b′ ∈ π^{-1}(R_b); fiber connectedness joins a to a′ within π^{-1}(R_a) and b′ to b within π^{-1}(R_b). Concatenation gives an admissible reversible path a → b, which is RC. ■

Remark 3.2 (the converse is ALP itself). The reverse implication RC ⟹ ALP does *not* hold in general and is deliberately omitted. RC supplies, for each pair a, b, *one* trajectory in $\mathcal{A}$ whose projection is *some* base path between R_a and R_b. ALP, by contrast, quantifies over *every* continuous path in $\mathscr{R}$, and an arbitrary γ from R_a to R_b need not be the projection of any admissible reversible trajectory. A connected total space over a connected base is not thereby a fibration — this is precisely the fact §7 rests on, that connectedness does not imply lifting. Asserting RC ⟹ ALP would either fail or silently reintroduce ALP as a hypothesis. The biconditional of the previous paper is therefore not warranted; only the forward implication above is, and it is all that Theorem 8.1 requires.

The "fiber connectedness" hypothesis is exactly the **phase-continuity tier** isolated previously: connectedness within a fiber is connectedness of phase motion at fixed carving. We carry it forward explicitly rather than absorbing it silently. $\mathscr{R}$-connectedness was reduced to capacity geometry in prior work and is treated here as established. One direction-of-quotient point must be confirmed, since §2 now takes $\mathscr{R}$ raw (pseudometric) with the floor-quotient $\mathscr{R}$/≈_$\mathscr{R}$ formed only later: connectedness pushes *forward* along the quotient map $\mathscr{R}$ → $\mathscr{R}$/≈_$\mathscr{R}$ but not backward, so what Proposition 3.1 consumes must be connectedness of the **raw** $\mathscr{R}$. If the prior capacity-geometry result is stated on the floor-quotient, it does not by itself deliver raw-$\mathscr{R}$ connectedness; we therefore take the established result to be raw-$\mathscr{R}$ connectedness (the stronger and safe-direction form), and note that this is the form the lifting argument requires. The remaining work is therefore ALP.

---

4. The Operational Ontology and the Sector Principle

We make explicit the foundational commitment the programme has assumed throughout, because the status of the Sector Principle depends entirely on it.

Axiom OO (Operational Ontology). A distinction is a structural feature of the substrate only insofar as it is operationally witnessable at the substrate's distinguishability scale. Distinctions with no possible operational witness at the floor are not features of the substrate; they are descriptive surplus.

OO is not introduced for this paper. It is the same commitment from which finite distinguishability, the commitment-event ontology, and the operational construction of capacity were drawn. We restate it only to fix the inference that follows.

Definition (operational witness). An operational witness for a distinction D is a procedure, composed of operations available at or above the distinguishability floor, whose floor-resolved outcome differs according to D.

Definition (sector boundary). A sector boundary in $\mathcal{A}$ is a partition of admissible reversible motion into classes such that motion is confined within a class — here, the putative wall separating phase motion from refinement motion.

Lemma 4.1 (Operational Sector Principle, OSP) [Conditional-on-OO].

A sector boundary that admits no operational witness is not a structural feature of $\mathcal{A}$.

Proof. A sector boundary is a distinction among admissible motions: it asserts that membership in one class versus another is a real fact about a motion. By OO, such a fact is a structural feature only if it admits an operational witness at the floor. A boundary with no witness therefore asserts a fact that OO does not recognize as structural. It is descriptive surplus and may be removed without altering any operationally accessible feature of $\mathcal{A}$. ■

This is the first reclassification. In the previous framing OSP was listed as a conjecture co-equal with FBI, which made the residue look like two open problems of comparable depth. Under OO, OSP is a lemma, not a conjecture. It is open only in the sense that OO itself is a foundational posit — but OO is *already assumed* everywhere in the programme, so admitting OSP as a separate conjecture would have double-counted the foundational commitment. The honest residue is therefore lighter than previously stated by exactly one conjecture.

OSP is the application of OO to one kind of object — a *distinction* (the sector boundary). OO applies equally to a second kind of object — a *function* on states (a probability assignment). That second application yields the Operational Indistinguishability Principle and, with it, an account of where ODG itself comes from. We develop this in §9–§12, after the sector argument is complete, and show in §12 that OSP and OIP are twin corollaries of a single principle.

---

5. Sub-Floor and Composite Identification

We now treat the second principle, and here finite distinguishability tightens rather than loosens the argument. The previous statement of FBI used "infinitesimal" phase and refinement displacements. There are no infinitesimals at a finite floor. We restate the identification at floor scale and, in doing so, split it into a part that is nearly automatic and a part that is the true open problem.

Definition. A displacement is *sub-floor* if its capacity displacement is below the distinguishability floor, and *supra-floor* otherwise.

Claim 5.1 (Sub-floor identification) [Conditional/near-trivial].

A sub-floor phase displacement and a sub-floor refinement displacement produce no distinguishable substrate states; the two are operationally identical.

Justification. A refinement displacement registers only through a change in the carving, and a sub-floor change in the carving is by definition not resolvable — it maps to "no operational change." A sub-floor phase displacement leaves the carving fixed and rotates phase by an unresolvable amount — again "no operational change." Both sub-floor motions therefore have the same floor-resolved description, namely the null one. The only assumption beyond the definition of the floor is that phase and refinement register *exclusively* through floor-resolvable capacity changes, which is the content of the operational construction of $\mathcal{A}$. ■

Claim 5.1 disposes of the local version of the identification almost for free. It is not, however, sufficient. A wall could be sub-floor-invisible step by step yet recoverable by *composing* many supra-floor steps — the integration, or holonomy, problem. Local indistinguishability does not entail global indistinguishability. This is precisely the gap that a continuum-infinitesimal phrasing concealed. The genuine residue is the composite statement:

Conjecture FBI-comp [Conjectural].

There is no protocol, built from finitely many floor-resolved operations, whose outcome distinguishes "this finite admissible path was phase motion" from "this finite admissible path was refinement motion." Phase and refinement directions are operationally indistinguishable not only sub-floor but for all finite composite paths.

FBI-comp is the physical core of the problem, now stated without appeal to infinitesimals and without smuggling in continuum structure that the substrate rejects. The remaining derivation problem for the programme is FBI-comp specifically — not a vague "FBI."

---

6. Operationally Invisible Sector Boundaries

We assemble the pieces.

Theorem 6.1 [Conditional-on-OO + FBI-comp].

Under OO, FBI-comp implies that any sector boundary separating phase from refinement motion is operationally invisible, hence not a structural feature of $\mathcal{A}$.

Proof. Suppose a sector boundary separates phase from refinement motion. An operational witness for this boundary would be a protocol of floor-resolved operations whose outcome differs according to which side of the boundary a motion lies on — equivalently, a protocol distinguishing phase from refinement motion. By Claim 5.1 no sub-floor protocol does so, and by FBI-comp no finite composite protocol does so. There is therefore no operational witness for the boundary. By OSP (Lemma 4.1), a witnessless boundary is not a structural feature of $\mathcal{A}$. ■

Note what this does and does not claim. It does not claim that "a fibered topology requires a sector boundary" — that direction was definitional unpacking in the prior draft and is dropped. It claims the converse and operative direction: granting FBI-comp, any such boundary is surplus and removable. The structure with the boundary removed we call **unsectored**.

A terminological caution. We avoid the word "fibered" for the boundary-bearing structure. In standard usage a fibration *possesses* the path-lifting property — the opposite of the obstruction we are describing. We use "sector-partitioned" for the boundary-bearing structure and "unsectored" for its removal, reserving "fiber" only for $\pi^{-1}(R)$, the phase content over a fixed carving.

Corollary 6.2 [Conditional-on-OO + FBI-comp].

$\mathcal{A}$ is unsectored: there is no admissible structural wall between phase and refinement motion.

---

7. From Unsectored Motion to the Lift Property

Here lies the step that the previous paper asserted and that does not hold for free. We isolate the missing premise rather than conceal it.

Removing the sector boundary makes $\mathcal{A}$ unsectored, but unsectoredness is a statement about the *absence of one particular partition*. Path-lifting (ALP) is a stronger property than the absence of a single wall: a continuous surjection over a connected total space can still fail to lift base paths, for reasons unrelated to any sector partition — failure of local triviality, holonomy obstructions, or preimages that thin out along a path. Connectedness does not imply lifting. The previous chain "single-topology ⇒ ALP" silently upgraded the one to the other.

We therefore name the premise that the chain actually needs.

Premise UO (Uniqueness of Obstruction) [Conditional].

The sector boundary of §4–6 is the sole obstruction to ALP. Equivalently: once $\mathcal{A}$ is unsectored, every continuous base path in $\mathscr{R}$ admits an admissible reversible lift.

UO is not yet proven. It is plausible in the present setting because the base topology of §2 was constructed from admissible transport, so each base step is by construction realizable in $\mathcal{A}$; the only structural feature that could block carrying such a path reversibly is a wall forbidding the carry — which is the sector boundary. But "the only feature we have named" is not "the only feature there is," and the holonomy possibility in particular must be excluded, not assumed away. We carry UO explicitly as a [Conditional] premise.

The §4–§7 consistency, made explicit. There is an apparent circle. §4 calls a witnessless wall *descriptive surplus* — not a structural feature at all — so on what grounds could it ever obstruct a lift, and why does UO treat "the sector boundary" as a candidate obstruction? The resolution turns on a single identification, which we elevate to a named result so that everything downstream can reference it explicitly rather than relying on a remark buried in discussion.

Bridge Theorem (obstruction ⟹ witness) [Conditional].

If a sector boundary obstructs ALP, then the obstruction is operationally witnessable.

Contrapositive form, which is what the §6→§7 chain actually consumes: an operationally unwitnessable sector boundary does not obstruct ALP.

Justification. A sector boundary obstructs ALP only by forbidding the continuation of some admissible reversible lift across it — i.e. by making the two sides inequivalent for the purpose of carrying a path. But an inequivalence that governs whether a motion may continue is, by the operational construction of admissibility, a distinction the substrate enacts: a procedure that attempts the crossing has a floor-resolved outcome (continuation permitted vs forbidden) that differs across the boundary. That differing outcome is an operational witness. Hence an obstructing boundary is witnessable. The status is [Conditional]: it depends on the identification of "forbids continuation" with "produces a floor-resolved difference," which is substantive, not definitional. ■

Remark (the converse is not claimed). Only obstructs ⟹ witnessable is asserted and used. The converse, witnessable ⟹ obstructs, is neither needed nor plausible — one can detect which side of a wall a motion is on without being forbidden to cross it, so a witnessable boundary need not obstruct lifting. Stating the theorem as the single implication the chain consumes avoids a gratuitous referee target.

Granting the Bridge Theorem, §4 and §7 are consistent and complementary. UO speaks counterfactually: *were* there a real wall, it would obstruct lifting — and by the theorem such an obstructing wall would be operationally witnessable. OO together with FBI-comp shows no operationally witnessable wall exists (§6); by the contrapositive there is therefore no obstructing wall, which is exactly the content UO needs on the sector side. So §4 removes the wall as surplus, and §7 records that the only thing that *could* have obstructed the lift was precisely such a wall — the same object under its two descriptions, "structural feature" and "lift obstruction," tied together by the Bridge Theorem. It is the hinge on which the apparent circularity between §4 and §7 turns into a single coherent argument.

Theorem 7.1 [Conditional-on-OO + FBI-comp + UO].

Under OO, FBI-comp, and UO,

$\mathcal{A}$ unsectored ⟹ ALP.

Proof. By Corollary 6.2, OO and FBI-comp give that $\mathcal{A}$ is unsectored. By UO, unsectoredness is exactly the condition under which every continuous base path lifts admissibly and reversibly, which is ALP. ■

7.1 Why UO Is Not an Arbitrary Extra Premise

UO is the most load-bearing remaining premise in the lifting argument, so it must not be hidden. We can nonetheless show it is far from arbitrary, by reducing the space of obstructions that could falsify it.

We are not considering an arbitrary topological projection. $\mathcal{R}$ was built from admissible transport (§2), and $\mathcal{A}$ is the space of admissible reversible pre-commitment configurations. A failure of ALP can therefore occur only if a base-admissible refinement path cannot be carried as an admissible reversible trajectory in $\mathcal{A}$. We organize the candidate failure modes into three types. Their **exhaustiveness is itself a [Conditional] claim**: standard lifting theory admits failure modes beyond these three — π failing to be open, or a global selection obstruction that is neither pure holonomy nor pure fiber-collapse — and the enumeration below should not be read as carrying tacit completeness. We treat the three as the dominant cases in the present construction and flag the residual possibility explicitly.

1. **Sector obstruction** — a wall separates phase and refinement motion into distinct admissibility sectors. This is exactly the boundary treated in §4–§6.
2. **Holonomy obstruction** — a lifted path returns with a nontrivial phase or structural mismatch.
3. **Degeneracy obstruction** — the fiber collapses, branches, or becomes ill-defined along the base path.

Holonomy is not an obstruction to lifting as such. It is an obstruction to path-independence or closure: a nontrivial holonomy means that different lifted paths around a loop return with different phase data, but it does not prevent a lift from existing — indeed, holonomy is *defined on* existing lifts around loops, so it presupposes their existence. Holonomy can therefore obstruct ALP only if it is upgraded into a sector wall that forbids continuation, i.e. only by becoming a case of obstruction type 1.

Degeneracy would require the fiber structure to fail along the base path. By construction every point of $\mathscr{R}$ is an admissible refinement, and every admissible refinement has at least one compatible substrate configuration, so $\pi^{-1}(R)$ is nonempty for every $R \in \mathscr{R}$. Pointwise nonemptiness eliminates the crudest degeneracy: the lift cannot fail simply for want of a target.

Lemma 7.1A (Obstruction Reduction) [Conditional].

Given the transport construction of $\mathscr{R}$ and the definition of $\mathcal{A}$ as admissible reversible pre-commitment configurations, failure of ALP requires a structural sector obstruction. Holonomy obstructs path-independence and degeneracy obstructs uniqueness, but neither eliminates the existence of a lift; only a structural prohibition on crossing between admissibility sectors does.

Status and caveat. The lemma is Conditional in three respects. First, it assumes that every admissible refinement carries a nonempty compatible fiber and that admissible transport paths meet no singular refinement at which π ceases to be defined — natural in the present construction, but to be checked against the full substrate formalism. Second, and more sharply: pointwise nonemptiness of fibers secures a lift *target at each point*, but ALP demands a **continuous** lift — a continuous selection along the whole path. A branching or collapsing fiber can leave every $\pi^{-1}(R)$ nonempty yet admit no continuous section across the degenerate locus. The honest reduction is therefore that degeneracy obstructs *continuity* of the lift rather than its mere pointwise existence. Lemma 7.1A is best read as reducing ALP failure to the union {sector obstruction} ∪ {continuity failure at a fiber degeneracy}. Excluding the latter — showing that admissible transport never forces a fiber degeneracy at which a continuous selection fails — is the residual content folded into UO alongside the sector boundary. Third, the reduction inherits the [Conditional] exhaustiveness of the three-type enumeration above: if a failure mode outside {sector, holonomy, degeneracy} is possible — an openness failure of π, or a global selection obstruction of neither pure type — it is not addressed here, and the union should be read as {sector obstruction} ∪ {continuity failure at a degeneracy} ∪ {residual modes}.

With this reduction, UO is no longer the bald assertion that the sector boundary is the only obstruction. It is the claim that the remaining apparent obstructions either presuppose lifting (holonomy) or, in the degenerate case, reduce to a continuity condition on admissible transport that the construction is expected to satisfy. UO's open content is thereby narrowed from "no obstruction of any kind" to "no fiber degeneracy of admissible transport breaks continuous selection" — a sharper and more checkable target.

---

8. Consequences for Reversible Connectedness

We combine §3 and §7.

Theorem 8.1 [Conditional-on-OO + FBI-comp + UO + phase continuity].

Under OO, FBI-comp, UO, $\mathscr{R}$-connectedness, and fiber (phase) connectedness,

RC holds.

Proof. Theorem 7.1 gives ALP. $\mathscr{R}$-connectedness is established by prior capacity-geometry work. Fiber connectedness is the phase-continuity tier. Proposition 3.1 then assembles $\mathscr{R}$-connectedness ∧ ALP ∧ fiber connectedness into RC. ■

The continuity that the phase axis already requires (fiber connectedness) does **not** automatically extend to refinement motion merely because the space is unsectored. It extends via ALP, which requires UO. This is the correction of substance over the previous draft: the transfer of continuity from phase to refinement is mediated by an explicit lifting premise, not by connectedness alone.

---

9. Operational Quotients from Finite Distinguishability

We now turn from the sector application of OO to its second application: probability. The argument in §9–§11 establishes that the geometry on which probability lives (ODG) and the principle constraining probability (OIP) share a single origin in finite distinguishability. The bridge is a quotient construction, and the construction must be built carefully, because the obvious version of it fails.

Let A denote the operational carrier — the space of substrate states equipped with whatever structure operational procedures can access. One is tempted to define an indistinguishability relation directly by pairwise resolvability:

$\psi \approx \varphi \Leftrightarrow$ no admissible operational procedure resolves ψ from φ.

This relation is reflexive and symmetric, but it is **not** transitive, and the failure is not incidental. It is the tolerance (sorites / Poincaré) structure: sub-floor differences accumulate. One may have $\psi \approx \varphi$ and $\varphi \approx \chi$ — each pair differing by less than the floor — while ψ and χ differ by nearly twice the floor and so are resolvable, giving $\psi \not\approx \chi$. A relation that is reflexive and symmetric but not transitive does **not** partition the carrier into classes. Quotienting by $\approx$ is therefore ill-posed, and any theorem that quotients "by operational indistinguishability" naively inherits this defect.

Finite distinguishability supplies the repair at the root of the hierarchy, through finite packing. But the repair needs more from finite packing than mere finiteness, and we state the stronger premise precisely.

Premise FP (Finite Packing) [Conditional, prior result]. At finite capacity the floor induces a **determinate, single-valued** resolution map

$r : A \to A\_op$

onto a finite operational state set A_op, surjective, assigning to each state — *including states near a resolution boundary* — a unique operational image. Equivalently, the floor induces a genuine finite partition of A, not merely the existence of finitely many distinct states.

The strengthening over "finitely many distinct states" is essential and is exactly where the sorites could otherwise re-enter. The entire $\approx \to \sim$ repair below works because r is a function: one projects each state *once* and compares images. If a borderline state sitting between two operational states had no determinate image — if assigning $r(\psi)$ were itself subject to the tolerance ambiguity — then r would not be well-defined as a function, and the same accumulation that defeats transitivity of $\approx$ would reappear in the supposed quotient. FP must therefore be read as asserting a clean partition with determinate boundary assignment, not the weaker counting claim. (The danger that the *image* of r is infinite or non-discrete is flagged separately in §14; the danger addressed here is the more delicate one that r fails to be a function at all.)

FP is not introduced here; it is the finite-packing result already established in the programme, carried forward as a premise in this stronger, determinate form. With it, define indistinguishability not pairwise but through the resolution map:

Definition (operational equivalence).

$\psi \sim \varphi \Leftrightarrow r(\psi) = r(\varphi)$.

Relation of the two quotients. This paper now carries two projections, and their relationship must be fixed to avoid the §2/§10 double-count and to keep the sector and probability sides commensurable. The refinement projection $\pi : \mathcal{A} \to \mathscr{R}$ sends a state to its carving, with fibers = phase. The resolution map $r : A \to A\_op$ floor-resolves *every* operational degree of freedom — both carving and phase. In general, then, A_op is the total set of a finite bundle over the floor-quotiented refinement space $\mathscr{R}/\approx\_\mathscr{R}$, with fiber the floor-quotiented phase content over each carving:

A_op = total set of a finite bundle over $\mathscr{R}/\approx\_\mathscr{R}$, fiber $\Phi\_R/\approx\_\Phi$ over [R].

Here r is the joint floor-resolution and π is, post-resolution, the induced projection onto the refinement factor: the operational refinement space $\mathscr{R}/\approx\_\mathscr{R}$ is the image of A_op under the induced projection. Thus π and r are not rivals — π names the carving coordinate and r floor-resolves all coordinates — and the floor-quotient that §10 applies to the pseudometric d of §2 is the refinement-factor part of the single quotient by r, performed once.

This bundle reduces to a **product**

$A\_op \cong (\mathscr{R}/\approx\_\mathscr{R}) \times (\Phi/\approx\_\Phi)$

only if the phase fibers Φ_R have constant operational cardinality over all carvings R — i.e. only if no fiber branches, collapses, or otherwise changes cardinality as R varies. We state this condition explicitly rather than assume it, because it is not free: a fiber that branches or collapses over different R is precisely the **degeneracy obstruction of §7.1**, which Lemma 7.1A and UO carry as *open* on the sector side. Were we to assert the product silently, the probability side would be granting exactly the no-degeneracy condition the sector side is still trying to discharge — a circularity.

We therefore make the dependence a feature, not a leak.

Shared non-degeneracy condition (ND). The phase fibers have constant operational cardinality over $\mathscr{R}/\approx_{\mathscr{R}}$. Equivalently: admissible transport meets no fiber degeneracy along any base path.

ND is the same structural fact under two descriptions. On the probability side it is what upgrades the bundle to the product factorization above. On the sector side it is the residual content of UO isolated in Lemma 7.1A — "no fiber degeneracy of admissible transport breaks continuous selection." So the factorization holding and (the degeneracy part of) UO holding are *one fact viewed from the two sides*. This is a sharper architectural claim than OO-sharing alone: the sector and probability programmes share not only the operational ontology but this specific assumption, ND, and it dovetails with the §14 remark that UO and FBI-comp are plausibly two faces of one fact. In what follows the product form is used only under ND, which is tagged [Conditional] wherever it is invoked.

Theorem 9.1 (Finite distinguishability generates an operational quotient) [Conditional on FP].

The relation ~ is an equivalence relation, and the operationally meaningful state space is the quotient carrier $A/\sim \,\cong A_{op}$.

Proof. As the kernel of the function r, the relation ~ is reflexive, symmetric, and transitive by construction — equality of images $r(\psi) = r(\varphi)$ inherits these properties from equality in A_{op}. Hence ~ partitions A into the fibers of r, and the quotient $A/\sim$ is in canonical bijection with the image A_{op}. Because r is, by FP, the finest distinction the substrate can operationally enact, operational structure cannot separate points within a fiber of r; it acts on the fibers, i.e. on $A/\sim$, not on individual representatives of A. ■

Two points deserve emphasis. First, the move from ≈ to ~ is exactly what defeats the sorites: rather than chaining pairwise "cannot-tell-apart" judgements (which accumulate), we project each state once to its operational image and then compare images (which does not). The resolution map is what discretizes the carrier. Second, the construction is grounded in FP and so sits *below* ODG in the hierarchy: it is finite packing, not ODG, that makes the quotient well-defined. This ordering matters in §12.

---

10. ODG as Quotient Geometry

We now identify operational distinguishability geometry with the geometry induced on A/~. The key is a single invariance lemma; once it is in hand, descent of every operational structure is immediate rather than gestural.

Lemma 10.1 (Operational definability implies ~-invariance up to resolution) [Conditional on FP].

Any structure on A that is defined operationally is constant on ~-classes *up to operational resolution*, and therefore descends to a well-defined structure on A/~.

Proof. Suppose a structure S took *operationally resolvable* different values on ψ and φ with ψ ~ φ. Reading off S is itself an operation at or above the floor; a resolvable difference in S would then be an operational procedure separating ψ from φ, contradicting r(ψ) = r(φ). Hence S cannot vary across a ~-class by any resolvable amount. For a discrete-valued operational structure this means exact constancy on fibers. For a continuous- or real-valued structure it means constancy *up to floor resolution*: sub-floor variation across a fiber is invisible but need not be exactly zero. In either case the floor-resolved value of S is constant on each fiber of r, so S descends to a well-defined structure on A/~. ■

The weaker "up to resolution" conclusion is the honest one and is exactly what the metric instance below exhibits: d is not literally zero across a ~-class but has *resolved* value zero, which is why its descent is realized as a Kolmogorov quotient rather than as exact pointwise vanishing.

Theorem 10.2 (ODG is quotient geometry) [Conditional on FP].

The carrier, sector structure, projection structure, transport structure, refinement structure, and finite-packing structure each descend to A/~, and the induced geometry is ODG. ODG is therefore not an independent postulate; it is the quotient geometry generated by finite distinguishability.

Proof. Each listed structure is, by its construction in the programme, defined operationally. By Lemma 10.1 each descends to A/~. The transport structure descends in a sharper sense: the pseudometric d of §2 has zero set exactly the floor-indistinguishable pairs, which under r is exactly ~, so d descends to an honest (non-degenerate) metric on A/~ — the Kolmogorov quotient of the pseudometric space. The collection of induced structures on A/~ is, by definition, ODG. ■

10.3 What the Quotient Result Does and Does Not Prove

The claim that ODG is quotient geometry must be read with precise scope. It does not say that ODG is derived from no assumptions. It rests on three prior ingredients:

1. finite packing supplies a finite operational state set A_op;

2. the resolution map $r : A \to A\_op$ identifies which substrate states are operationally identical;
3. the structures entering ODG are operationally defined.

Given those ingredients, the quotient result is forced. The theorem therefore has the form:

Finite Packing + Operational Construction ⇒ Operational Quotient ⇒ ODG as the induced geometry on the quotient.

The result is not that ODG appears from nothing; it is that, once finite packing and operational construction are granted, ODG cannot carry information below the operational quotient. Any proposed ODG distinction separating two states with the same image under r would be extra structure not visible to the substrate — and, by Lemma 10.1, no operationally defined structure can draw such a distinction.

This is the theorem's real content. ODG is not an independent postulate in the sense of adding a separate state space over and above finite distinguishability; it is the geometry of the quotient that finite distinguishability generates. Overstated as "ODG is not postulated at all," the claim would invite the objection that the conclusion was placed in the definitions; stated with the three ingredients explicit, it is airtight.

ODG nonetheless remains conditional on the operational construction of its constituents. A critic who rejects that construction is not compelled by the quotient theorem; the theorem shows only that *within* the VERSF operational framework ODG is not an additional primitive. Its epistemic status is therefore [Conditional on FP + operational construction], the same footing as Theorem 10.2.

---

11. The Operational Indistinguishability Principle

We can now treat probability. Let μ be a probability assignment on states.

Definition (operationally meaningful assignment). μ is operationally meaningful if it factors through operational structure — equivalently, if its value depends on a state only through what operational procedures can access about that state.

Theorem 11.1 (Necessity of OIP) [Conditional-on-OO, given FP].

Every operationally meaningful probability assignment is constant on ~-classes and therefore factors uniquely through A/~. This constancy is exactly the Operational Indistinguishability Principle.

Proof. Let $\psi \sim \varphi$, and suppose $\mu(\psi) \neq \mu(\varphi)$. Then μ assigns different values to states that operational structure cannot separate; by the definition of operational meaningfulness, μ is therefore not operationally meaningful. Contrapositively, an operationally meaningful μ is constant on ~-classes. By the universal property of the quotient, such a μ factors uniquely as $\mu = \bar{\mu} \circ r$ for a unique $\bar{\mu}$ on $A/\sim \;\cong A_{op}$. Constancy on ~-classes is the statement of OIP. ■

Two clarifications keep this honest. First, the step "μ distinguishes operationally identical states ⇒ μ is not operationally meaningful" is an *identity*, not an inference: it merely restates the definition of operational meaningfulness. The theorem's substantive content lies entirely in (i) Theorem 9.1's quotient, which rests on FP, and (ii) the premise that probability *must* be operationally meaningful. The latter premise is where OO carries the load: the natural objection — that probability might be a substrate-fundamental, operationally hidden object (a ψ-ontic or hidden-variable quantity) legitimately separating operationally identical states — is rejected precisely by OO, which denies structural status to anything with no operational witness. This is why OIP earns the tag [Conditional-on-OO]: it is no stronger and no weaker than OO itself.

Second, the unique factor $\bar{\mu}$ lives on the finite set A_{op}. The forced domain of probability is therefore exactly the finite set of operationally distinct outcomes — the correct domain for a Born rule over distinguishable results, with no overreach beyond it.

---

12. OSP and OIP as Twin Corollaries

We can now state the architectural result that motivates the preceding sections.

Proposition 12.1 (Twin corollaries of OO) [Conditional-on-OO].

OSP and OIP are the single principle OO applied to two different objects:

— Applied to a *distinction* (a sector boundary): a distinction with no operational witness is not substrate structure. This is OSP (Lemma 4.1): a witnessless wall is not a feature of $\mathcal{A}$.

— Applied to a *function on states* (a probability assignment): a function that separates states carrying no operational distinction is not operationally meaningful. This is OIP (Theorem 11.1): μ cannot distinguish what the substrate cannot distinguish.

Proof. Both are instances of the OO schema "X is admissible only if the distinction X draws has an operational witness," with X ranging over sector boundaries in the first case and over level sets of probability assignments in the second. In each case the absence of a witness is supplied operationally — by FBI-comp for the sector boundary (§6) and by Theorem 9.1's quotient for the probability assignment (§11) — and OO then removes the witnessless structure. ■

Fraternal, not identical. The symmetry is at the level of *principle*, not of *standing*, and the two must not be read as equally secure. Both apply OO to an object, but the witness-absence each relies on has sharply different epistemic status. On the sector side, the absence of a witness is delivered by FBI-comp, which is [Conjectural] and remains the principal open problem. On the probability side, it is delivered by Theorem 9.1's quotient, which is a *theorem* given FP. OSP-via-FBI-comp and OIP-via-Theorem-9.1 are therefore fraternal corollaries — same parent OO, different current footing: one waits on a conjecture, the other is discharged modulo a prior result. The shared schema should not be allowed to disguise this asymmetry of standing.

The unifying slogan is one sentence: **reality contains only what the substrate could in principle detect.** Pointed at walls it gives OSP; pointed at probabilities it gives OIP. Two principles that earlier drafts of the programme carried as independent postulates collapse into one foundational commitment with two faces.

Strengthened conceptual hierarchy. The probability chain previously read

ODG → OIP → ODG-Compatibility → UMP → Born Rule,

with ODG and OIP both entering as independent inputs. The quotient construction gives ODG and OIP a common origin and lengthens the chain at its root, while §12.1 adds the second input the UMP step actually requires:

Finite Packing Geometry ↠ Operational Quotient Structure ↠ ODG ⊨ OIP ⟹ ODG-Compatibility ⟹ UMP ⟹ Born Rule,

↑

RC

with RC feeding the UMP step as its connectivity precondition (§12.1). The arrows are deliberately of three distinct kinds, and conflating them would overstate the result:

— ↠ *generative*: the earlier object constructs the later one (finite packing builds the quotient via r; the quotient carries ODG as its induced geometry).

— ⊨ *constraining*: ODG, as quotient geometry, forces OIP as a property any operationally meaningful assignment must satisfy.

— ⟹ *derivational*: the standard downstream implications, unchanged from prior work.

— ↑ *precondition*: RC supplies a hypothesis (connectedness of the admissible space) that UMP consumes.

The Born endpoint inherits the weaker of its two branches. The left-to-right chain is not "discharged modulo FP," because the UMP step also consumes RC, and RC depends through Theorem 8.1 on FBI-comp, which is [Conjectural]. The Born-rule endpoint therefore inherits the *weaker* of the standings of its two input branches: discharged-modulo-FP on the OIP branch, but only conjectural on the RC branch, so conjectural overall pending FBI-comp. This must not be confused with the "fraternal, not identical" point of §12: that asymmetry concerns OSP versus OIP as *principles*, whereas the Born rule depends on *both* branches and so cannot be more secure than the RC branch. The gain over the previous hierarchy is structural, not a change in the endpoint's ultimate standing: ODG and OIP no longer enter as two separate assumptions but descend from the single root of finite distinguishability, exactly as OSP and OIP descend from the single root OO — and the conjectural residue is now correctly localized in the RC branch rather than diffused across the chain.

12.1 Locating RC Relative to the Probability Chain

The paper now carries two threads that each bear on the Born rule, and their relationship must be stated or the reader cannot tell whether there is one target or two. The abstract and §1 frame Reversible Connectedness (RC) as the Born-rule residue on the sector side; §9–§12 give the chain ODG ⊨ OIP ⟹ ODG-Compatibility ⟹ UMP ⟹ Born Rule on the probability side. These are **not two independent derivations of two Born rules**, nor are they yet proven to be one fully unified derivation. The coherent reading, which we adopt, is that they address *different sub-problems of a single target*:

— The probability chain supplies the **measure**: given the operational state space, OIP forces any operationally meaningful assignment to factor through A_op, and the uniqueness argument (UMP) fixes that assignment as the Born measure.

— RC supplies the **space on which that measure lives**: the uniqueness-of-measure argument requires that the admissible reversible motion space be connected (and the relevant continuity hold), so that a single measure is determined across the whole space rather than independently on disconnected sectors. A failure of RC would leave the UMP step quantifying over a fragmented space, on which "the" unique measure need not be globally defined.

On this reading RC is a **connectivity precondition the probability chain consumes**, not a parallel route: the chain's UMP step has an implicit hypothesis — connectedness of the admissible space — that RC is exactly what discharges. The two threads are therefore sequential, RC ⟶ (precondition for) UMP, joining at the measure-uniqueness step.

Status of this reconciliation [Conditional]. The precise form of the connectivity hypothesis that UMP requires is fixed in the UMP paper, not here; what is asserted is the direction of dependence (RC feeds UMP), not the exact strength of the hypothesis. If the UMP argument turns out to require less than full RC — say, connectedness of A_op rather than of $\mathcal{A}$ — the dependence stands but its quantitative form weakens. We flag this so the single-target picture is not mistaken for a completed unification.

---

13. The Residue, Restated

The remaining burden is now precise. Already discharged or reduced to prior results: $\mathscr{R}$-connectedness; the sub-floor identification (Claim 5.1); the reclassification of OSP as a consequence of OO (Lemma 4.1); the obstruction reduction (Lemma 7.1A); and the common derivation of ODG and OIP from finite packing (§9–§12). What remains is a finite list of named checks, split by side.

The sector side requires:

1. **FBI-comp** [Conjectural] — no finite floor-resolved protocol distinguishes phase motion from refinement motion.
2. **UO / Obstruction Reduction** [Conditional] — failure of ALP requires a genuine sector obstruction, not merely holonomy or non-unique lifting; by Lemma 7.1A the open content narrows to a continuity condition, that no fiber degeneracy of admissible transport breaks continuous selection of a lift.
3. **Phase continuity** [Conditional] — the fiber-connectedness tier inherited from the phase-axis argument.

The probability side requires:

1. **FP** [Conditional, prior result] — finite packing yields a finite operational state set A_op via a determinate, single-valued resolution map.
2. **OO** [Axiom, programme-wide] — operationally unwitnessable distinctions are not substrate structure.
3. **Operational construction of ODG's constituents** [Conditional] — so that the quotient-geometry theorem (10.2) applies.
4. **ND (non-degeneracy)** [Conditional] — the phase fibers have constant operational cardinality, so the §9 bundle reduces to the product factorization. This is *not* an independent probability-side assumption: ND is the same fact as the degeneracy content of UO (Lemma 7.1A), so the factorization holding on the probability side and UO holding on the sector side are one condition viewed from two sides.

The programme's residue is therefore no longer a broad question about probability. It is the finite list above, and it is shorter than its length suggests, because ND (item 4) coincides with the degeneracy part of UO (sector item 2). No probability-side item is a new free-standing conjecture: FP and OO are standing commitments, the constituent-construction requirement is the same one ODG already carries, and ND is shared with the sector side.

The deepest remaining technical question is whether FBI-comp and UO can be derived from finite distinguishability and the construction of the admissible-motion space — and, as noted in §14, whether they are in fact one fact viewed from two sides.

---

14. Limitations

This paper does not prove FBI-comp. Claim 5.1 settles only the sub-floor case; the composite case is open.

This paper does not prove UO. The holonomy possibility — that finite composites of supra-floor moves reconstruct a lift obstruction absent at the level of single steps — is not excluded here, and is the same structural risk that FBI-comp addresses on the identification side. UO and FBI-comp are plausibly two faces of one fact; establishing that they coincide would itself be a result. Lemma 7.1A (§7.1) narrows UO's burden — holonomy presupposes lifts, and pointwise-nonempty fibers exclude crude degeneracy — but does not discharge it: the residual obligation is to show that admissible transport never forces a fiber degeneracy at which a continuous lift fails to exist.

This paper does not prove phase continuity. It is carried as the fiber-connectedness hypothesis of Proposition 3.1.

OSP is no stronger than OO. If the operational ontology were weakened, OSP would lose its [Conditional-on-OO] standing and revert to an independent posit.

The quotient construction (§9) rests on FP, and on FP in its *strong* form. Two distinct dangers must be excluded. The coarser: if finite packing did not deliver a genuinely finite operational state set — if r had infinite or non-discrete image — Theorem 9.1 would not yield a well-posed partition. The subtler, and more dangerous: if r failed to be well-defined as a function — if a borderline state near a resolution boundary had no determinate image — then the tolerance ambiguity that defeats transitivity of ≈ would reappear inside the supposed quotient, and the ≈ → ~ repair would collapse. FP must therefore assert a determinate, single-valued resolution including on boundary states, not merely "finitely many distinct states." The repair survives only if the floor induces a determinate r; both the discreteness of the image and the determinacy of the assignment are load-bearing and are carried as premises, not proven here.

Theorem 10.2 (ODG as quotient geometry) is near-tautological given that ODG's constituents are operationally defined; its content is that prior operational construction, not the descent step. We do not claim ODG arises from no commitments, only that it carries no information beyond A/~.

OIP (§11) is no stronger than OO, for the same reason as OSP. In particular, a ψ-ontic or hidden-variable reading of probability — one that grants structural status to operationally hidden state separations — is not refuted within this paper; it is excluded by OO, which such a reading rejects. Against an opponent who denies OO, OIP is not established.

---

15. Conclusion

The remaining obstacle to Reversible Connectedness is no longer connectivity, geometry, probability, or individuation, and it is no longer a pair of co-equal conjectures.

Under the operational ontology already assumed throughout the programme, a wall between phase and refinement motion is admissible only if something at the distinguishability floor could detect it. The sub-floor case shows nothing can detect it locally. The whole question is therefore whether anything can detect it by composition — whether finitely many floor-resolved operations, chained together, recover a distinction that no single step provides.

If finite distinguishability forbids that — FBI-comp — and if the sector boundary is the only thing that ever blocked a lift — UO — then phase and refinement motion belong to one continuous admissible motion space, and the Born-rule residue is discharged.

The question is now sharply localized and stated in the substrate's own terms:

Can finitely many floor-resolved distinctions, composed, witness a wall that no single distinction can see?

Beyond the sector question, the same operational standard reorganizes the probability programme. Finite distinguishability, through finite packing, generates an operational quotient of the state carrier; the geometry of that quotient is ODG; and the principle that probability cannot distinguish operationally identical states is then a theorem about quotient spaces rather than a postulate. The witnessless-wall principle (OSP) and the indistinguishable-probability principle (OIP) are one principle — reality contains only what the substrate could in principle detect — applied to a distinction and to a function respectively. The sector programme and the probability programme thereby share a single foundation, and the open burden in each is correctly located: FBI-comp and UO on the sector side, and the operational ontology together with finite packing on the probability side.