Whether the Carvings Flow

Why Reversible Connectedness Is the Same *Species* of Condition as the Phase-Axis Continuity the Programme Already Owes — and Why Whether It Is the Same *Theorem* Turns on a Single Question About the Topology of Admissible Motion

Keith Taylor VERSF Theoretical Physics Programme versf-eos.com

---

For the General Reader

Picture reality, before anything has happened, as a region holding a fixed amount of "possibility" — a capacity. Nothing in it is yet labelled; the labels (the outcomes) are *carved* out of the region's own capacity when reality finally commits to one way of cutting it up. The companion papers established two things about these carvings. They are made from within — no outcome is handed in from outside (the Internality Axiom). And the weight each carving carries is also fixed from within, by the capacity it holds, not by the company it keeps (the previous paper, on contextual weight).

That left exactly one question between this picture and the squared-amplitude rule of quantum mechanics. It is not about what reality *contains* — that ledger is now settled. It is about whether reality can *move*: whether one way of carving the region can flow continuously into another, or whether the carvings sit in sealed boxes with no path between them. We have called the "yes" answer **Reversible Connectedness** — RC. It is the last gate.

Now, the founding rule is silent about motion. It hands over a region and a capacity; it does not say which motions are allowed. So RC looks, at first, like a genuinely new thing we would have to assume — a fact about dynamics the rule cannot give us. This paper asks whether that is really so, or whether we already owe the very same continuity somewhere else and simply have not noticed it is the same.

And we do owe it elsewhere. To get the squared-amplitude rule at all, the programme leans on a quieter motion being continuous: the smooth rotation of a possibility's *phase* while its carving stays fixed. Phase-spin is continuous; that is assumed, or proven, on what we call the phase axis. RC asks whether *carving-change* is continuous on the refinement axis. Two motions, both allowed, both reversible, both before commitment. Why should one flow smoothly and the other not?

The honest answer this paper reaches is in three parts. First, the two motions are the *same kind of thing* — there is no difference in their credentials, so RC is not a new *species* of assumption; it is the carving-axis cousin of a continuity we already lean on. Second, when we open RC up, one half of it turns out to be cheap: that the carvings form one connected landscape (rather than scattered islands) follows almost automatically once we measure distance between carvings by how much capacity must be shifted to get from one to another — there is no real difficulty there, and saying so lets us see where the difficulty actually is. What is left of RC — that the connected landscape can actually be *travelled* by allowed motion — is the genuine remainder, and it is the exact twin of phase-continuity. Third, and this is the catch we refuse to hide: showing two motions are the same kind of thing does not show that proving one proves the other. There is a real gap — the same gap that, in the previous paper, stopped "the books balance under spinning" from delivering "the books balance under re-carving." Whether that gap closes is the one question still open, and we leave it open rather than wish it shut.

So the last gate is no longer a strange new lock; it is a familiar kind of lock, with one bolt whose mechanism we can now describe exactly. But we are careful about what that buys. We have not made the difficulty smaller — travelling the landscape is exactly as hard a question as RC ever was. We have made it *single and sharp*: the whole of it now sits in one precise question about whether allowed motion can carry one carving into another, a question of the same kind the programme already faces for phase-spin. That is a real gain in clarity, not a reduction in difficulty — and we say so plainly rather than let the image of a "nearly-turned lock" suggest otherwise. Whether the carvings flow now waits on that single bolt.

Abstract

The Born-exponent arc stands at **ℓ^2 ⟸ IA ∧ Obstruction B**, with IA proven, Obstruction B affirmative ⟺ **RC** (Reversible Connectedness), and RC the sole open gate — a continuity premise on the *dynamics* axis the founding rule FP1 cannot adjudicate. The previous companion, closing the data axis, flagged RC twice as "plausibly not new either, the refinement-axis kin of the phase-axis continuity upgrade the programme already owes." This paper opens that parenthetical. The goal is *not* to prove the Born rule; it is to determine whether RC is a genuinely new premise or already implied by the phase-axis continuity the programme leans on elsewhere.

We separate two theses with sharply different premises:

- **Thesis A (category):** RC and phase continuity are the same *species* of condition — both assert that admissible, reversible, pre-commitment substrate motion is continuous along connected paths. If so, RC is not a new *kind* of dynamical debt.
- **Thesis B (closure):** the arc closes, because that shared continuity is itself *established* — which requires the phase-side continuity to be proven (not merely owed) and to *transfer* to the refinement sector.

The central results:

1. **Same-Species classification — SS (proven, modulo inherited admissibility criteria).** A substrate motion is admissible iff capacity-preserving, reversible iff invertible, pre-commitment iff prior to the selection of a refinement. Phase rotation and refinement-change satisfy all three. So both belong to one class of admissible reversible pre-commitment motion (curves in the motion space $\mathcal{A}$, §2.4). SS is a *classification*, hence FP1-adjudicable in the same way IA was — it concerns what the substrate admits, on the data axis. **Thesis A follows from SS.**
2. **The RC decomposition, in a constructed projection (proven).** We equip the carving space $\mathscr{R}$ with a **transport metric** (distance = capacity displacement, the minimal topology generated by Vol_op), build the motion space $\mathcal{A}$ of (carving, phase) configurations with projection $\pi : \mathcal{A} \to \mathscr{R}$, and define admissible motion as a *restricted* curve class (capacity-preserving, reversible — a proper subclass of continuous curves). Then RC ⇔ **$\mathscr{R}$-connectedness** (the base is path-connected) ∧ **ALP, the Admissible Lift Property** (every base path admits a lift that is an admissible reversible motion). RC is thereby a *lifting problem* about π, not a vague continuity claim.
3. **$\mathscr{R}$-connectedness is cheap, not a theorem — and that is the clarification.** Under the transport metric, base path-connectedness is near-automatic from capacity-continuity (modulo an asserted, not proved, judgement that Vol_op carries no internal sector-invariant). We *demote* it rather than dress it as a hard result: the transport topology is built from capacity alone, precisely so connectedness is not engineered into the open sets, and the demotion concentrates the entire difficulty in ALP, where it belongs. Earlier framing overstated this half; the honest statement is that the programme's difficulty never lived here.
4. **ALP is phase continuity's twin (same species, by SS); whether it is the same *theorem* is a controllability property of the admissible-direction distribution (open).** ALP and phase continuity are the same predicate — admits-an-admissible-reversible-lift / is-an-admissible-reversible-motion — on base paths versus fiber paths. Same species (SS), but base-lift is **cross-refinement** and fiber-motion **intra-refinement** — and the previous paper's lesson was that an intra-refinement fact does not reach a cross-refinement fact for free. The transfer holds iff the distribution $D \subset T\mathcal{A}$ of admissible directions is **bracket-generating** (Chow–Rashevskii: admissible motion connects across fibers) rather than involutive-and-fiber-tangent (foliated into the fibers; no base-transverse admissible direction; ALP fails). The obstruction lives in D — *not* in the point-set topology of $\mathcal{A}$ (we use "single-topology" only as informal shorthand for "D bracket-generating," and correct the tempting holonomy framing in §10.5). We derive bracket-generation from an **Operational Sector Principle (OSP** — any admissibility-*sector* boundary needs an operational witness, proven modulo the programme's operational individuation, held selectively to sector boundaries not intra-class distinctions), together with **6.3A-bridge** (the motion-directions are the complete witness-set for a fiber-tangent involutive segregation of D — conjectural) and **6.3B** (the fiber and base directions are operationally indistinguishable at the distinguishability floor — conjectural). An application of OSP *shifts the burden*: the fibered horn must exhibit an operational witness, independent of the motions, that D is fiber-tangent involutive — which on the programme's commitments finite distinguishability forbids. We nonetheless label the bridge and 6.3B conjectural and the stratification conjectural-with-toy-coherence, and *refuse to assume* bracket-generation — asserting it directly would simply *be* RC renamed (§7). (Because admissible motion is

reversible, ALP may demand bracket-generation *with a measure-preserving lift*, strictly stronger than controllability — §10.5.)

Net, at the honest tiers. **Proven:** RC is the same *kind* of condition as phase continuity (Thesis A, via SS); RC ⟺ ALP, with $\mathscr{R}$-connectedness an entailed cheap precondition (ALP ⟹ $\mathscr{R}$-connectedness), so ALP carries the whole content; ALP is PhC's twin; OSP holds modulo the programme's operationalism. **Cheap, demoted:** $\mathscr{R}$-connectedness, near-automatic under the transport topology (modulo one asserted no-invariant judgement). **Conjectural, the locus of the remaining debate:** 6.3A-bridge and 6.3B, and the stratification reconciling the sharp-fibered regime (where the twinning is stated) with the blurred-at-floor regime (where the lift transfers) — supported by a toy model proving coherence, not actuality. **Open (Thesis B):** closure requires (i) the phase-side continuity *proven* not posited (a tier we flag), and (ii) the admissible-direction distribution bracket-generating (loosely, π "single-topology"), which we leave conjectural and do not force. So the defensible statement is: **RC is localized, not lessened.** It is not new in *kind*; its base half is cheap; its dynamical half is one named lift-property of a constructed projection, the same species as a phase-axis property already owed. The difficulty is concentrated, undiminished, in that single lift-property — which is a real gain in precision, not a reduction in difficulty, and we are careful not to let the rhetoric outrun that. The arc closes iff phase continuity is proven and lifts across the gap — and that lift-property, not RC as a monolith, is what the squaring now waits on. (This is an internal-consistency result; whether the framework tracks physics is a separate question it does not address.)

Contents

Scope and Conditional Status

This is the sixth companion to *The Squaring Residue*, after the linearity paper, *Bath or Ledger*, *No Pre-Individuation and the Seal-Trichotomy*, and *No Contextual Weight*, and inherits their framework and conventions. It does not assert the bath reading by preference. It opens the one object the previous papers left as the sole open gate — RC, the dynamical continuity premise whose affirmative resolution is Obstruction B — and determines whether RC is a new premise or already implied by the phase-axis continuity the programme owes elsewhere.

The result is that RC is not new in *kind*, and is no longer a monolith. We separate a **category** claim (RC is the same species as phase continuity — Thesis A, reached) from a **closure** claim (the arc closes — Thesis B, conditional on the phase tier and on one open structural question). The cleanest way to see the structure is to construct the projection $\pi : \mathcal{A} \to \mathcal{R}$ (the carving space under a transport metric, the motion space, and the projection forgetting phase) and decompose RC = $\mathcal{R}$-connectedness ∧ ALP (the Admissible Lift Property). $\mathcal{R}$-connectedness is then *cheap* — near-automatic under the transport topology, demoted from the substantial theorem earlier drafts claimed. ALP is phase continuity's twin. What is genuinely open is whether the admissible-direction distribution is bracket-generating (loosely, whether π is "single-topology"), so that phase continuity *lifts* from fiber paths to base paths — the fiber–base gap, located in the admissible directions (not the point-set topology of $\mathcal{A}$, §10.5), described precisely as a controllability property and declined as an assumption. We are explicit that this is *localization*, not reduction: the lift-property is exactly as hard as RC was, now single and sharp.

Epistemic labelling is maintained throughout:

- **proven** — follows from the stated inputs (FP1 as stated; the inherited companion results IA, the normalization axioms, the Diagonal-Torus selection and Possibility-Connectivity Theorem; the §2.3 Declaration of *No Pre-Individuation*) by an argument given here;
- **conditional** — follows given a named, separately-open premise or a companion statement to be confirmed;
- **conjectural** — asserted as plausible with structural support but no argument approaching proof.

SS (Same-Species) is **proven** modulo the inherited admissibility criteria (§5). RC ⟺ ALP, with $\mathcal{R}$-connectedness an entailed cheap precondition (ALP ⟹ $\mathcal{R}$-connectedness), is **proven** (§3). $\mathcal{R}$-connectedness is **cheap** — near-automatic under the transport topology of §2.4, modulo an asserted (not proved) judgement that Vol_op carries no internal sector-invariant (Lemma 6.0, Remark 6.0.1); we demote it rather than dress it as a theorem. ALP ≡ PhC-for-refinement is **proven** via SS (§6.2). OSP (the Operational Sector Principle, §6.3) is **proven modulo the programme's operational individuation**, held selectively to sector boundaries. Bracket-generation of the admissible-direction distribution (loosely, π "single-topology") — whether base paths admissibly lift — is **conjectural**, derived from OSP ∧ 6.3A-bridge (the motion-directions are the complete witness-set for a fiber-tangent involutive segregation of D) ∧ 6.3B (the motion-directions are operationally indistinguishable at the floor), with the burden shifted onto the fibered horn but the gap explicitly not assumed (§6.3, §7); a strictly-stronger second tier — bracket-generation *with a measure-preserving lift*, should "reversible" demand it — is left open (§10.5). The stratification reconciling the sharp and blurred regimes is **conjectural-with-toy-coherence** (Remark 6.3.1). The tier of the phase-side continuity (proven vs posited) is flagged **conditional** (§4); it decides Thesis B's reachability. RC is **open** in exactly the lift-property these isolate, and we say so.

A methodological note, carried because it is the same discipline. The previous papers refused to let a contested conclusion ride in on a word, and refused the symmetric trap of proving a structural fact and then announcing the bath. *No Contextual Weight* added a third: it declined the "conservation forces additivity" shortcut, because the inherited conservation is *intra*-refinement and the thing wanted was *cross*-refinement. That exact distinction is load-bearing here, against *us*: the temptation this paper must resist is to let intra-refinement phase continuity flow, unargued, into the cross-refinement Admissible Lift Property — the same intra/cross leak, now in our own favour. We name the gap and leave it open rather than wish it shut.

Time is emergent throughout; "pre-factual" / "pre-commitment" denote the substrate configuration before an irreversible commitment event. A commitment event selects one admissible refinement of the unresolved region (the Declaration, *No Pre-Individuation* §2.3). "Phase axis" denotes amplitude/phase motion at fixed carving; "refinement axis" denotes change of carving.

---

1. Introduction: the last gate

Across five companions the squared-amplitude rule was chased down to a single residue and the residue was anatomised until one conjunct remained. *Bath or Ledger* localised the mixing leg to a fork in the conservation structure of pre-factual weight and tied it to the *Packing* paper's refinement obstruction, Obstruction B. *No Pre-Individuation* disproved the primitive ledger, proved the Internality Axiom IA, and showed Obstruction B's affirmative resolution equivalent (given IA) to **Reversible Connectedness (RC)** — whether the admissible reversible substrate motion connects the carvings — a continuity question on the *dynamics* axis FP1 cannot adjudicate, because FP1 legislates capacity and is silent on motion. *No Contextual Weight* then closed the data axis: the decomposition-independence principle PC, long flagged as "the one genuinely new bookkeeping commitment," was shown to be IA on the weight axis (universality) plus inherited measure-additivity (composition), not a new *kind* of debt. That paper's closing ledger read:

$\ell^2 \Leftarrow$ IA $\wedge$ Obstruction B, IA proven, Obstruction B affirmative $\Leftrightarrow$ RC, RC open.

and it twice let drop a parenthetical about the lone survivor — that RC is "plausibly not new either, being the refinement-axis kin of the phase-axis continuity upgrade the programme already owes." That parenthetical is doing precisely what PC's old caption — *the one genuinely new bookkeeping commitment, to be stated and checked* — did before the previous paper opened it. It is a flagged box, not an opened one. This paper opens it.

The goal is narrow and we state it to forestall the obvious misreading. **We do not here prove the Born rule, and we do not prove RC.** We determine RC's *status*: whether it is a genuinely new dynamical premise, or whether the programme already owes the same continuity on the phase axis and RC is that continuity wearing refinement clothes. The distinction matters because the two outcomes leave the programme in very different places. If RC is new, the squaring rests on a dynamical posit FP1 cannot reach, full stop. If RC is the phase continuity's kin, then the squaring rests on *one* continuity, owed once, instantiated twice — and the open question shifts from "is RC true?" to "does the continuity we already lean on cover both instances?"

The temptation, having one open conjunct that is dynamical, is to treat RC as a single irreducible dynamical fact and ask only whether to assume it. We show that picture is too coarse in two ways. First, RC is not monolithically dynamical: opened up in a constructed projection $\pi : \mathcal{A} \to \mathscr{R}$, one half of it — that the carvings form a connected base — is, under the natural transport topology on the carving space, a near-automatic fact with no motion in it. Second, the dynamical remainder is not a new continuity but the twin of one already in play on the phase axis. What is genuinely open, after these two moves, is neither "is RC true" nor "is RC new in kind" — it is whether the continuity proven (or posited) on the phase axis *lifts* to the refinement axis across a gap the previous paper's own discipline warns may not be free to cross. We are explicit that this is *localization*, not reduction: the lifting question is exactly as hard as RC was, but it is now single, named, and sharp.

So the contribution is to relocate and localize RC — not to shrink it (the open question is as hard as it was, §10.7), but to isolate it, by argument and not relabelling, as the one structural question on which closure now turns. We are at pains, throughout, to apply against ourselves the exact lesson *No Contextual Weight* applied to the conservation shortcut: an intra-refinement fact does not reach a cross-refinement claim merely by resembling it.

---

2. Inherited setup: two motions, one substrate

2.1 The carrier, capacity, and the conserved normalization

From the companions we take as given the carrier H0: the pre-factual state $\psi = (c_1, \ldots, c\_d) \in \mathbb{C}^d$, components a path-sum read off relative to a chosen refinement of the unresolved region M, and the conserved separable normalization $N(\psi) = \Sigma_i\, h(|c_i|)$ with axioms N1 (BCB-additivity), C2 (diagonal phase invariance), N4 (faithfulness), H-mono. The capacity Vol_op assigns to each admissible carving its capacity content, with **Capacity Additivity** $\mathrm{Vol\_op}(M) = \Sigma_i\, \mathrm{Vol\_op}(M_i)$ (the Declaration, *No Pre-Individuation* §2.3). Two inherited theorems fix the phase-side picture: the **Diagonal-Torus selection** (continuous off-diagonal mixing $\Leftrightarrow$ ℓ^2) and the **Possibility-Connectivity Theorem** (mixing $\Leftrightarrow$ bath).

2.2 Two motions

The substrate, before commitment, admits motion of two apparent kinds, and the whole paper is about their relation.

Phase motion (intra-refinement). At a *fixed* carving R, the amplitudes may rotate: the off-diagonal mixing of the Diagonal-Torus selection, a continuous action that leaves N invariant. The carving is held; the phases move. This is motion on the **phase axis**.

Refinement motion (cross-refinement). The carving itself may change: M is re-carved from one admissible refinement R into another R′. The capacity is held (Vol_op(M) is fixed; CA holds in every carving); the partition moves. This is motion on the **refinement axis**.

The first is the motion phase continuity is *about*. The second is the motion RC is *about*. They share three credentials — admissibility, reversibility, pre-commitment — which §5 makes precise and which is the whole basis of the same-species claim. They differ in what they hold fixed and what they move: phase motion holds the carving and moves within it; refinement motion moves the carving. Whether that difference is a difference of *kind* or merely of *direction* in one motion space is the question of §6.

2.3 What "the programme owes" means, and why RC sits beside it

FP1 supplies capacity; it does not supply a dynamics. So *every* statement that a particular substrate motion is admissible-and-continuous is, strictly, a premise FP1 does not deliver — it is *owed*. The programme already runs on one such premise: the continuity of phase motion, leaned on wherever the Diagonal-Torus selection is invoked (§4 examines its tier). RC is a second statement of the same shape — that refinement motion is admissible-and-continuous-enough to connect the carvings. The previous paper's parenthetical is the conjecture that these two owed premises are *one*. This paper tests it.

2.4 The spaces, made explicit: $\mathscr{R}$, $\mathcal{A}$, π, and admissible motion

The companion papers supply Vol_op, admissible refinements, and distinguishability capacity, but — as far as this sequence has developed — *no formal topology on the refinement space*. The load-bearing claims below (connected, path, lift) are topological, so we must specify the structure rather than argue by analogy to a fiber bundle that has not been built. We do so minimally, and flag every choice.

Definition (the carving space $\mathscr{R}$ and its transport metric). $\mathscr{R}$ is the set of admissible refinements (partitions of Vol_op(M)). Equip it with the **optimal-transport (Wasserstein-type) metric**

d(R, R′) = inf over couplings of { capacity-content cost of a coupling between R and R′ },

where a *coupling* is a static transport plan — a joint apportionment of Vol_op(M) whose two marginals are the partitions R and R′ — and the infimum is over plans, **not** over time-parametrized trajectories. The induced metric topology is the topology on $\mathscr{R}$ throughout.

Disclosure (a choice, minimal — and the firewall it protects). This metric is not inherited from prior work; we adopt it as the minimal topology *generated by the primitive object Vol_op* — distance is the optimal capacity-content coupling cost, nothing else. It is essential that the metric is defined by *couplings* (configuration-level objects: joint apportionments), **not** by transformations or processes carrying R to R′. A coupling is static; no parametrized motion enters the metric. This is what keeps the configuration/motion firewall standing: "$\mathscr{R}$-connectedness is motion-free" is true precisely because the base distance is a coupling cost, not a transformation cost. Were d defined by an infimum over re-apportionment *processes*, the base metric would be motion-defined and the firewall would collapse — so we are explicit that it is not. Alternative topologies are possible, but introducing additional admissibility structure into the topology itself would risk relocating the dynamical question into the definition of the space; the coupling metric refuses that, built from capacity alone, so any difficulty that survives is a difficulty about *motion* (ALP), not about a topology chosen to produce it.

Definition (the motion space $\mathcal{A}$ and the projection π). $\mathcal{A}$ is the space of admissible reversible pre-commitment substrate configurations: pairs (carving, phase-configuration), carrying the carving's position in $\mathscr{R}$ together with the amplitude/phase data the carrier H0 assigns at that carving. The projection

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

forgets the phase data, returning the carving. Fibers $\pi^{-1}(R)$ are the phase-configurations over a fixed carving R; the base is $\mathscr{R}$. We give $\mathcal{A}$ the metric topology from the product of the transport metric on $\mathscr{R}$ and a phase metric on the fibers. (We do *not* assume π is a fiber bundle or that $\mathcal{A}$ is a global product; whether it is is precisely the open question of §6.)

Definition (admissible motion as a restricted curve class). An **admissible reversible motion** is a continuous curve in $\mathcal{A}$ satisfying the admissibility criteria of Lemma 5.2 — capacity-preserving, reversible, pre-commitment. This is a *proper subclass* of continuous curves: not every continuous path in $\mathcal{A}$ is an admissible motion. The distinction is the whole hinge of what follows — a continuous curve through configurations need not be an admissible motion — and it is the configuration/motion (point/path) distinction of §6.1 in the language of the constructed spaces.

These four definitions are this paper's, adopted minimally and flagged as such; a programme-canonical topology on $\mathscr{R}$ or specification of π and the motion class, should one be fixed later, could redistribute the load between connectedness and lifting (§6), which is itself the honest disclosure the construction demands.

---

3. RC stated precisely, as a lifting problem

RC entered the arc as the affirmative resolution of Obstruction B, proven equivalent to it given IA. It was never stated as an object in its own right; we do that now, in the constructed spaces of

§2.4, and decompose it — at which point it resolves into a connectedness claim about $\mathscr{R}$ and a *lifting* claim about π.

RC (Reversible Connectedness). For any two admissible refinements R, R′ of M, there is an admissible reversible motion carrying R to R′ — equivalently, an admissible reversible curve in $\mathcal{A}$ whose projection under π runs from R to R′.

Two things must hold, and in the constructed spaces they are cleanly distinct — one about the base $\mathscr{R}$, one about lifts through π:

$\mathscr{R}$-connectedness (the base claim). The carving space $\mathscr{R}$, in the transport topology, is path-connected: between any R, R′ there is a continuous path $\gamma : [0,1] \to \mathscr{R}$ with $\gamma(0) = R$, $\gamma(1) = R'$. This is a property of the base alone — whether the carvings form one landscape or metric-separated islands — with no reference to motion in $\mathcal{A}$.

Admissible Lift Property (ALP — the lifting claim). For every pair R, R′ there is *some* continuous path γ in $\mathscr{R}$ from R to R′ admitting an **admissible reversible lift**: a curve $\lambda : [0,1] \to \mathcal{A}$ that is an admissible reversible motion (§2.4) with $\pi \circ \lambda = \gamma$. The bite is *admissible reversible*, not merely continuous — a continuous lift may exist freely while no *admissible* lift does. (We state ALP in the existence-of-a-liftable-path form RC requires; the stronger "every path lifts" is not needed and not claimed.) ALP is the property of π on which everything now turns; it replaces the looser "Traversability" of earlier framing with the precise lifting question.

Proposition 3.1 (RC ⇔ ALP, with $\mathscr{R}$-connectedness an entailed precondition — proven). RC ⇔ ALP. Moreover ALP ⟹ $\mathscr{R}$-connectedness, so equivalently RC ⇔ $\mathscr{R}$-connectedness ∧ ALP; but the conjunction is not a decomposition into *independent* halves — $\mathscr{R}$-connectedness is entailed by ALP and is factored out only to show the difficulty does not lie in it.

Proof. (ALP ⟹ $\mathscr{R}$-connectedness): a path admitting an admissible lift is in particular a path, so ALP's liftable connecting path witnesses path-connectedness of the base for each pair. (RC ⇔ ALP): (⟸) ALP supplies, for each R, R′, a connecting path γ with an admissible reversible lift λ; then λ is an admissible reversible motion whose projection carries R to R′ — RC. (⟹) If RC holds, for each R, R′ there is an admissible reversible motion λ carrying R to R′; its projection $\pi \circ \lambda$ is a connecting path admitting the lift λ — ALP for that pair. Quantifying over R, R′ gives the equivalence; $\mathscr{R}$-connectedness follows from either side as the entailed precondition. ∎

So, strictly, **RC ⇔ ALP** — ALP carries everything. We nonetheless surface $\mathscr{R}$-connectedness as a named precondition, not because it does independent work (it does not — it is entailed by ALP), but because factoring it out and showing it *cheap* (§6.1) localises where the difficulty is *not*: not in whether the carvings form a connected base, but in whether admissible motion lifts the connecting paths. This is the lever of the paper. RC as a monolith fused a base fact with a lifting fact; once separated, the base fact is cheap and entailed, and the lifting fact (ALP) is the whole gate.

Remark 3.1.1 (this is the move PC underwent, on the refinement axis). *No Contextual Weight* took PC — flagged as one monolithic new commitment — and split it into a universality strand that was IA in disguise and a composition strand that was inherited bookkeeping. We take RC and factor it as ALP with its cheap entailed precondition ($\mathscr{R}$-connectedness, near-automatic under the transport topology), exhibiting that the whole content is the lifting claim, which is phase continuity's kin. The discipline is the same: do not argue about the monolith; locate the content precisely — here, in ALP, with the base precondition shown not to be where the difficulty lives.

Remark 3.1.2 (why this is not just renaming Traversability). Recasting the dynamical half as ALP is not cosmetic. "Traversability" invited the reading that, once a continuous family of admissible carvings exists ($\mathscr{R}$-connected), motion along it is near-free — the near-vacuity worry. The lifting formulation blocks that reading by construction: ALP asks for a lift *that is an admissible reversible motion*, and admissible motions are a proper subclass of continuous curves (§2.4). A path can exist downstairs (in $\mathscr{R}$) with no admissible lift upstairs (in $\mathscr{A}$) — that is exactly a connected base with no admissible section over it, the rigorous form of "connected but untraversable." So $\mathscr{R}$-connectedness does not pre-empt ALP; the gap between them is the gap between a continuous curve and an admissible motion, which is the fiber–base gap stated as a property of π.

---

4. Phase continuity stated precisely — and its tier

The reduction's payoff depends on what phase continuity *is* and, decisively for Thesis B, on whether it is **proven** or **posited**. We state it and mark the fork.

Phase continuity (PhC). The off-diagonal mixing at fixed carving — the Diagonal-Torus action on amplitudes — is a *continuous* one-parameter (torus) action: the substrate moves through phase configurations along continuous paths, reversibly, before commitment.

The Diagonal-Torus selection invokes PhC to pass from the admissible mixing to ℓ^2: it is the continuity of the torus action that the functional-equation/rotation argument consumes. The question that decides Thesis B is whether that continuity is *derived* within the selection or *built into* the choice of a torus.

- **PhC posited.** If the Diagonal-Torus selection *chooses* a continuous (torus, U(1)-type) mixing group as the admissible structure — continuity entering as part of the selection's setup rather than as a theorem — then PhC is itself *owed*, not proven. In that case the best this paper can reach is **Thesis A**: RC is the same species as an owed premise, which consolidates two owed premises into one species but proves neither. Closure (Thesis B) is then off the table on the phase side, independently of anything about the refinement axis.
- **PhC derived.** If the continuity of the mixing is *derived* — from a deeper continuity of the substrate (e.g. from finite distinguishability forcing a continuum of admissible phase

configurations, or from a proven connectedness of the mixing group) — then PhC is proven, and Thesis B becomes reachable *provided* the fiber–base gap (§6) closes.

We do not here adjudicate the fork; it is a fact about the linearity companion's Diagonal-Torus construction, and the companion is authoritative. We flag it **conditional** and as the *first* thing to confirm, exactly parallel to the previous paper's N4-tier check. **The honest reading of the programme's current state is that PhC is at least leaned on, and may be posited rather than proven** — in which case this paper's deliverable is firmly Thesis A, and Thesis B is a target conditional on upgrading PhC. We write the paper to be robust to the fork: SS, the decomposition, and $\mathscr{R}$-connectedness hold regardless; only the *closure* claim depends on PhC's tier.

Proposition 4.1 (No closure bypasses the phase tier — proven). There is no route to Thesis B (closure) that avoids the admissible-lift property of motion being *proven* somewhere. In particular, a direct proof of ALP is a proof of PhC-for-refinement, so it settles the phase tier rather than sidestepping it.

Proof. Suppose, toward bypassing the tier, one proves ALP directly without leaning on PhC. By Theorem 6.2, ALP is the predicate "admits an admissible reversible lift / is an admissible reversible motion" applied to base paths — exactly PhC's predicate applied to the refinement sector. A proof of ALP is therefore a proof that admissible reversible motion lifts the base paths: PhC-for-refinement, proven. So the supposed bypass *is* a proof of the lift-property at issue, on the refinement sector; it settles the tier (in the "derived" direction, for that sector) rather than avoiding it. Conversely, closure requires ALP (RC ⟺ ALP, §3), and ALP just shown to be a lift-property of admissible motion; so closure requires that property proven somewhere — either on the fiber sector and transferred (§6.3), or on the base sector directly. Either way it is proven, not posited. Hence no closure route leaves the tier unfaced. ■

This is why the phase tier is genuinely load-bearing rather than one option among several: every path to closure runs through *proving* the continuity of admissible motion, and the only question is on which sector and whether it transfers.

---

5. Same-species: both motions are one class (Thesis A)

We now make precise the sense in which phase motion and refinement motion are the *same kind of thing*, and show the claim is a *classification* — adjudicated by the substrate's own admissibility criteria, hence on the data axis, hence within FP1's reach in the way IA was.

The substrate's criteria for a pre-commitment motion to be admissible reversible motion are inherited:

- **Admissible** ⟺ capacity-preserving: the motion leaves the total capacity content Vol_op(M) invariant (it redistributes or re-presents capacity but neither creates nor destroys it).
- **Reversible** ⟺ invertible: the motion has an admissible inverse.
- **Pre-commitment** ⟺ prior to selection: the motion acts on the unresolved substrate, before the commitment event selects a refinement.

Lemma 5.1 (Phase motion ∈ class — proven). Phase motion is admissible, reversible, pre-commitment. Admissible: the off-diagonal mixing leaves N invariant (the inherited conservation of the Diagonal-Torus action), and N faithfully tracks capacity (N4), so capacity content is preserved. Reversible: the torus action is invertible (rotate back). Pre-commitment: it acts at fixed unresolved carving, before selection. ■

Lemma 5.2 (Refinement motion ∈ class — proven, modulo CA). Refinement motion is admissible, reversible, pre-commitment. Admissible: by Capacity Additivity, $\mathrm{Vol\_op}(M) = \Sigma_i \mathrm{Vol\_op}(M_i)$ in every admissible carving, so re-carving holds the total capacity content fixed — capacity is preserved across the change of partition. Reversible: a re-carving $R \rightsquigarrow R'$ has the inverse re-carving $R' \rightsquigarrow R$ (both admissible). Pre-commitment: re-carving acts on the unresolved region, before selection collapses to one carving. ■

Theorem 5.3 (Same-Species, SS — proven modulo inherited criteria). Phase motion and refinement motion belong to one class of admissible reversible pre-commitment motion — the admissible curves in the motion space $\mathcal{A}$ (§2.4). They share all three admissibility credentials and differ in none *of them*. (They do differ in one respect — fiber-vs-base locus — which §6 shows is the dynamically decisive axis; SS's claim is precisely that this difference is *not* a difference in admissibility-class membership.)

Proof. Immediate from Lemmas 5.1, 5.2: both satisfy admissibility, reversibility, pre-commitment — the full set of *credentials* defining membership in $\mathcal{A}$. No admissibility criterion distinguishes them. They are distinguished by fiber-vs-base locus (what each holds fixed, §5.6), but that is not a membership criterion; it is a difference *within* the class, and §6 is devoted to whether it is dynamically consequential. ■

Corollary 5.4 (Thesis A — proven via SS). RC is not a new *kind* of dynamical debt. RC's content (by §3) is the lift-property (ALP) of refinement motion; phase continuity's content is the continuity of phase motion; by SS these are lift/continuity conditions on *one admissible-motion class* (the admissible curves in $\mathcal{A}$). So RC is, in *kind* — in admissibility-class — the same condition as PhC: both assert that motion in that class is continuous along connected paths. RC is the refinement-axis instance; PhC the phase-axis instance. (That the two instances may yet come apart — the base-lift instance not following from the fiber instance — is the fiber/base locus difference of §5.6, decisive in §6; "same kind" is the class claim, not a claim that the instances are interderivable.) ■

Remark 5.5 (SS is a classification, hence data-axis, hence FP1-eligible). SS does not assert that either motion *is* continuous — that is §6's separate question. It asserts that they are the same *kind* of motion, by the substrate's admissibility criteria. This is a claim about what the substrate *admits* — the data axis — and it is settled the way IA was settled: by reading off the substrate's own individuation of admissible motion. We are entitled to it on the same ground the previous papers were entitled to IA. What we are *not* yet entitled to is the inference from "same kind" to "same continuity"; that is the fiber–base gap, and it is dynamical, and §6 keeps it firmly separate from SS.

Remark 5.6 (answering Question 2: what makes refinement motion different?). By SS, *nothing in its credentials*. Phase and refinement motion are not distinguished as admissible reversible pre-commitment motions. The only difference is what each holds fixed: phase motion fixes the carving and moves within it (motion in a *fiber*); refinement motion moves the carving (motion in the *base*). Whether that fiber-versus-base difference is a difference of dynamical *kind* — two sectors that must be argued separately — or a mere difference of *direction* within one motion space, is precisely §6. The disanalysis of the two motions, if there is one, is located not in the motion but in the *topology of the motion space*.

---

6. The fiber–base gap: a lift-property of π, not a metaphor

SS (§5) and the decomposition (§3) have done their work; now comes the hard part, and we make it the centre of the paper rather than glide over it. The danger is exactly the one *No Contextual Weight* identified and refused — letting an intra-refinement fact flow unargued into a cross-refinement claim — and here the leak would be *in our favour*, which makes it more tempting and more important to block. Throughout, the question is a property of the constructed projection $\pi : \mathcal{A} \to \mathscr{R}$ (§2.4): whether base paths admit admissible reversible lifts.

6.1 $\mathscr{R}$-connectedness is near-automatic under the transport topology — and that is a clarification, not a theorem

In earlier framing we treated $\mathscr{R}$-connectedness as a substantial theorem. With the transport topology of §2.4 written down, it is not — and we now say so plainly, because the honesty sharpens the paper rather than weakening it. Under a metric whose distance *is* capacity displacement, the carving space is connected almost by construction once capacity is a connected continuum; the programme's real difficulty never lived here. Recognising that is the point of demoting the claim: it concentrates the entire burden where it belongs, on the lifting question (ALP).

We keep the configuration/motion distinction explicit, because it is what guarantees the demotion does *not* leak into ALP.

The configuration/motion (point/path) distinction. "Admissible" has two senses (§2.4). An *admissible carving* is a configuration — a capacity-preserving partition of Vol_op(M), a *point* of

$\mathscr{R}$. An *admissible motion* is a curve satisfying Lemma 5.2 — a *path* in $\mathcal{A}$, and a proper subclass of continuous curves. $\mathscr{R}$-connectedness quantifies over points (is there a continuous path of carvings); ALP quantifies over paths (does an admissible *motion* lift it). The first is settled by the transport topology; the second is not settled by it at all.

Lemma 6.0 (Connectedness reduces to capacity-continuity and the absence of a sector-invariant — near-definitional under the transport metric). In the transport topology, $\mathscr{R}$ fails to be path-connected only if some admissibility-separating invariant partitions it into metric-separated components — and by IA that invariant is either (1) external to Vol_op(M) or (2) an internal conserved structural invariant of Vol_op(M).

Proof. In the transport metric, d(R, R′) is the optimal capacity-content coupling cost between the partitions R and R′ (§2.4). Consider the interpolation $\{R_t\}$ of dividing-boundary positions sliding through the capacity continuum from R (t=0) to R′ (t=1). This is continuous in the transport metric: a small displacement of boundary positions admits a coupling moving only the capacity-content between the old and new boundaries, whose cost → 0 as the displacement → 0, so d(R_s, R_t) → 0 as s → t — small boundary displacement gives small coupling cost. (Note $\{R_t\}$ is a family of *configurations* indexed by t, not a parametrized motion; the metric and the path are configuration-level throughout.) The interpolation is a path in *$\mathscr{R}$ provided every R_t is an admissible carving*. By IA, admissibility-as-a-configuration is a predicate on partitions of Vol_op(M). If no admissibility predicate separates R from R′, every R_t is admissible and $\{R_t\}$ is a path — connectedness. So disconnection requires an admissibility predicate taking different values across components: a separating invariant, by IA of kind (1) or (2). ■

Remark 6.0.1 (the lemma is closer to definitional than "theorem" suggests — and we do not pretend otherwise). Once IA is granted — admissibility fixed by Vol_op-data and nothing else — "no separating invariant ⇒ connected" is nearly immediate; the lemma largely re-expresses that admissibility is invariant-determined, now in the transport topology where the interpolating path is automatically continuous. We flag this honestly: the lemma is a clarification of *where* an obstruction could sit, not a hard-won connectedness result. Its real content is the **conditional judgement that Vol_op carries no internal sector-invariant** (kind 2; kind 1 is already excluded by the prior papers). And that judgement is **asserted, not proved**: "capacity is a single additive content (CA)" is a plausibility note, not a demonstration that CA admits no conserved refinement-separating structure. We state it at exactly that tier.

Proposition 6.1 ($\mathscr{R}$-connectedness — proven under the transport topology, modulo the no-internal-invariant judgement). If Vol_op is a connected continuous measure and carries no internal refinement-separating invariant, then $\mathscr{R}$ is path-connected in the transport topology. No dynamical premise is used; every object is a configuration (a point of $\mathscr{R}$), no motion appears.

Proof. Immediate from Lemma 6.0: kind (1) excluded by the prior papers (*No Pre-Individuation*; *No Contextual Weight* §9), kind (2) excluded by hypothesis, so no separating invariant exists and the interpolation $\{R_t\}$ is a path of admissible carvings. ■

Remark 6.1.1 (the demotion is the clarification — and it does not touch ALP). Two consequences, both honest. First, $\mathscr{R}$-connectedness is *cheap*: near-automatic under the transport topology, resting on capacity-continuity (inherited) and the no-internal-invariant judgement (asserted at plausibility level, §10.3). It is not the locus of difficulty and we no longer present it as one. Second — and this is why the demotion costs nothing downstream — cheap connectedness does *not* shrink ALP. $\mathscr{R}$-connectedness gives a continuous path of *configurations* $\{R_t\}$; ALP asks for an admissible *motion* lifting it. In a fibered $\mathcal{A}$ the base can be perfectly path-connected (the $\{R_t\}$ exist) while *no admissible reversible motion lifts the path* — a connected base with no admissible section over it. So the gap between connectedness and ALP is the gap between a continuous curve downstairs and an admissible lift upstairs; demoting the first leaves the second exactly where it was, undiminished. The decomposition delivers a cheap base half and a full lifting residue, with no double-counting and no overlap.

So, post-§6.1, RC = ($\mathscr{R}$-connectedness: cheap, base-level, near-automatic under the transport topology) ∧ (ALP: the lifting residue — does an admissible reversible motion lift the connected base paths). The remainder of §6 is about ALP, which carries the entire difficulty.

6.2 ALP is phase continuity's twin — the same predicate on fiber paths and base paths

ALP asks whether base paths admit admissible reversible lifts. Phase continuity asks whether fiber paths *are* admissible reversible motions. In the constructed spaces these are visibly the same predicate on two families of paths through $\mathcal{A}$.

Theorem 6.2 (ALP ≡ PhC-for-refinement, by SS — proven, same predicate on two domains). "Admits an admissible reversible lift / is an admissible reversible motion" is one predicate on continuous paths in $\mathcal{A}$ (by SS, §5, the motion class is one class). Phase continuity applies it to *fiber* paths (fixed-carving phase motion is admissible reversible motion); ALP applies it to lifts of *base* paths (refinement motion along the connected $\{R_t\}$ is admissible reversible motion). Same predicate, two domains of paths — the same *species*, not yet the same *statement*: a predicate's holding on fiber paths does not entail its holding on base-path lifts, which is precisely §6.3. (Stating the twinning needs the fiber/base distinction *sharp*; §6.3's transfer needs it to *dissolve* — reconciled in Remark 6.3.1.)

Proof. PhC (§4): fiber paths (continuous torus paths at fixed carving) are admissible reversible motions. ALP (§3): base paths $\{R_t\}$ admit lifts λ that are admissible reversible motions. Both predicate "is / admits an admissible reversible motion" of continuous paths in $\mathcal{A}$; by SS the predicate ranges over one motion class. One predicate, two path-families. ■

So ALP is not a *new* continuity premise. It is the phase continuity the programme already owes (or has proven), the same predicate asked of base-path lifts instead of fiber paths — the precise content of the previous paper's parenthetical, now exhibited in the constructed π. **ALP is PhC's twin: same species, different domain.** Thesis A is fully in hand. Whether same-species becomes same-statement — whether the fiber instance *delivers* the base-lift instance — is §6.3. (One caution carried forward: because the predicate is *admissible reversible* motion, ALP inherits whatever measure-preservation "reversible" demands; mere connectivity of base to fiber may not suffice, and the gate may be strictly stronger than bracket-generation — §10.5.)

6.3 The gap: fiber-traversal does not entail base-traversal for free

Here we apply the previous paper's lesson against ourselves, and it bites.

PhC is **intra-refinement**: it concerns motion *within* a fixed carving (the fiber). ALP is **cross-refinement**: it concerns lifting a path *between* carvings (the base). *No Contextual Weight* established — as its cleanest result — that an intra-refinement fact (dynamical conservation, invariance under the within-carving reversible maps) does *not* reach a cross-refinement fact (subdivision-invariance) for free; the inference "intra, therefore cross" was an equivocation that, disambiguated, either smuggled the cross-claim or restated it. The structurally identical inference threatens here: "PhC (intra, fiber), therefore ALP (cross, base lift)" would, unargued, be the same leak — now flowing toward the conclusion we want, which is exactly when discipline is hardest and most necessary.

Whether the inference is legitimate turns on the topology of $\mathcal{A}$:

The fiber–base structure, and the right object: the admissible-direction distribution. $\pi : \mathcal{A} \to \mathscr{R}$ projects (carving, phase) to carving; fibers are phase-configurations over a fixed carving, the base is $\mathscr{R}$. Because admissible motion is a *restricted curve class* (§2.4) — not all continuous curves — the object that governs ALP is not the point-set topology of $\mathcal{A}$ but the **distribution D $\subset$ T$\mathcal{A}$ of admissible directions**: the directions a curve may take and remain an admissible motion. ALP — base paths admit admissible lifts — is then a *controllability* question for D, and the dichotomy is sharp and standard.

Two possibilities, and by **Chow–Rashevskii** they give opposite verdicts on whether admissible motion connects across fibers:

- **D involutive and fiber-tangent (the fibered horn).** If D is involutive — closed under Lie bracket — and tangent to the fibers, then by Frobenius it integrates to a foliation by the phase-fibers: every admissible direction lies *within* a fiber, no admissible motion has any base component, and base paths admit no admissible lift. ALP fails. Phase continuity is then the lift property of fiber paths only; it does not reach base paths, exactly as intra- and cross-refinement conservation were two independent facts last paper. **PhC does not transfer; ALP is a genuinely separate premise; Thesis B fails.** This is the disanalysis horn of Question 2, now precise: the two motions are the same *species* (both in the admissible class) yet base-lift fails because the class's *directions* do not reach off the fibers.
- **D bracket-generating (the connecting horn).** If D is bracket-generating — iterated Lie brackets of admissible directions span T$\mathcal{A}$, including base-transverse directions — then by Chow–Rashevskii admissible motion connects any two points of $\mathcal{A}$, base paths admit admissible lifts, and ALP holds. PhC and ALP are then two instances of one fact: D reaches everywhere. **PhC (if proven) delivers ALP; Thesis B is reachable.**

The open gate, correctly located. Is the admissible-direction distribution D *involutive and fiber-tangent* (foliated into the fibers — ALP fails) or *bracket-generating* (base-transverse — ALP holds)? This is the gate, and it lives in D, **not** in the point-set topology of $\mathcal{A}$. We use "single-topology" only as informal shorthand for "D bracket-generating"; the obstruction is a property of the admissible-direction distribution, which is why same-species (both motions in the class) is compatible with base-lift failure (the class's directions involutive, tangent to fibers).

This relocation matters and corrects an earlier framing (see §10.5): the gap is not whether $\mathcal{A}$ is "one space or two" as point-sets, but whether the *restricted* admissible directions bracket-generate across the fibers. We give the substrate-level reason to expect bracket-generation, splitting it — for the sharpest referee target — into a near-definitional admissibility principle and a specific physical claim about D at the floor.

Proposition 6.3.0 (Operational Sector Principle, OSP — proven modulo the programme's operational individuation). Any admissibility-*sector* distinction — any putative boundary partitioning the admissible-motion class into sectors — possesses an operational witness, or it is not a substrate-level distinction.

Proof. By the operational individuation of admissible structure that licensed IA and the external-charge exclusions, the admissibility facts are exactly the facts with operational content: admissible structure is individuated by what the substrate can register, not by structure laid over it. A purported sector boundary with no operational witness is a partition of the admissible class into pieces that differ in nothing the substrate registers — a distinction in no operational fact — hence not an admissibility fact. ■

The selectivity guard (why OSP does not over-collapse). OSP is a verificationist principle and such principles characteristically over-collapse — dissolving any distinction lacking an operational footprint, including ones the programme wants to keep. We block the cascade by restricting OSP's scope precisely: it applies to *sector boundaries* (partitions of the admissible class into sectors), **not** to distinctions *within* a sector. Two physically different but both-admissible configurations are distinguished *by the configurations themselves* — an intra-class distinction with an obvious operational witness (they differ) — and OSP says nothing against it. OSP bites only a putative boundary that would split the admissible class while having no witness. The fiber/base boundary is exactly such a candidate sector boundary (it would partition admissible motion into a fiber sector and a base sector); the distinctions the programme keeps are intra-class. We flag the residual risk honestly: the line between "sector boundary" and "intra-class distinction" is not perfectly sharp, and a critic who recategorises the fiber/base boundary as intra-class (or a kept distinction as a sector) would move the debate — but that recategorisation, not an unrestricted verificationism, is what they would have to argue.

With OSP in hand, 6.3A splits into a near-definitional core and a conjectural bridge — which is where the real debate now lives.

6.3A-core (from OSP — near-definitional). A fiber/base sector boundary with no operational witness is not a substrate-level distinction. A *fiber-tangent involutive* D — admissible directions confined to the fibers — is exactly such a sector boundary: it would partition admissible motion into "within-fiber, admissible" and "base-transverse, inadmissible." So D bracket-generates *unless* the involutive-fiber-tangent partition has an operational witness. (OSP applied to the one boundary at issue; carries OSP's tier.)

6.3A-bridge (the residual conjecture — completeness of motion-witnesses). The operational distinguishability of the two motion-directions is the *only* candidate witness of the fiber/base partition of D: there is no operational witness, independent of the motions, that some directions are "base-transverse and inadmissible" while others are "within-fiber and admissible." Equivalently, the motion-directions are a complete witness-set for the would-be sector boundary in D. **Conjectural**, and the precise locus of debate: a critic accepting OSP and 6.3B can resist bracket-generation only by exhibiting an *independent* operational witness that D is fiber-tangent involutive.

Conjecture 6.3B (Fiber–Base Identification — conjectural, the specific physical claim). Phase and refinement motion-*directions* are operationally indistinguishable at the distinguishability floor: below resolution there is no fact as to whether a sub-resolution admissible direction is a phase displacement within a carving or a boundary displacement between carvings — the fiber and base *directions* in D are not operationally separable at the floor. (Indistinguishability of *directions*, not of executed motions — §7.)

We must be careful translating this into the distribution language, because the translation crosses exactly the line §7's second-depth guard polices, and the distribution formalism makes the crossing more visible than the configuration language did. Distinguish two readings:

- **6.3B-config (purely configurational — what §7 permits).** Two sub-resolution *displaced configurations* — one obtained by a phase-displacement, one by a boundary-displacement — are operationally indistinguishable at the floor. This asserts nothing about tangent admissible directions; it compares configurations.
- **6.3B-dist (the distribution claim — D is bracket-generating at the floor).** The admissible-direction distribution D has base-transverse directions whose brackets reach off the fibers.

These are **not** the same claim, and the gap is precisely the second-depth circularity: 6.3B-dist already speaks of *admissible base-transverse tangent directions existing*, which is the existence-of-refinement-motion the second-depth guard (§7) forbids assuming. So the step **6.3B-config ⟹ 6.3B-dist is itself a derivation the eventual proof owes** — it must *derive* the existence of admissible base-transverse tangent directions from configuration-indistinguishability, not read it off the distribution. We flag this explicitly: where we write "6.3B's content is that D is bracket-generating," that is the *target* (6.3B-dist), reached from the *premise* (6.3B-config) only via a derivation that does not presuppose base-transverse admissible directions. §7 states the guard; here we mark that the config→distribution translation is where it bites.

The derivation, in distribution terms. From the three:

OSP ∧ 6.3A-bridge ∧ 6.3B ⟹ D bracket-generating ⟹ (by Chow–Rashevskii, admissible motion connects fibers; PhC delivers admissible base-lifts = ALP) ⟹ (with $\mathscr{R}$-connectedness, entailed) RC.

By 6.3B the fiber and base directions are operationally indistinguishable at the floor; by 6.3A-bridge that indistinguishability is the only candidate witness of a fiber/base partition of D, so the partition has no witness; by OSP (6.3A-core) a witnessless sector boundary is not a substrate distinction — so D carries no involutive fiber-tangent segregation, i.e. D is bracket-generating. Then by Chow–Rashevskii admissible motion connects across fibers, base paths admit admissible lifts (ALP), and with $\mathscr{R}$-connectedness (entailed, §3) RC follows.

The burden-shift (an application of OSP + 6.3A-core, not an independent argument). Granting OSP and 6.3B's indistinguishability, any fiber-tangent involutive segregation of D becomes operationally witnessless, hence — by OSP — not a substrate distinction. The burden then shifts onto the fibered horn: it must produce **an operational witness, independent of the motions, that D is fiber-tangent involutive** (an independent registrable structure segregating admissible directions into a within-fiber foliation) — exactly what 6.3A-bridge conjectures does not exist, and exactly the witnessless structure OSP forbids. So the fibered horn owes a concrete construction — an independent witness of involutivity — which on the programme's operational commitments is what finite distinguishability is built to forbid. This does not *close* the gap (6.3A-bridge and 6.3B remain conjectural; we do not assume RC), but it reverses the default: bracket-generation becomes the position that costs nothing, fiber-tangent involutivity the one owing an independent witness.

Remark 6.3.1 (the stratification, downgraded from mechanism to conjecture-with-toy-model). A referee will press a real tension. The theorems of §5–6.2 *need* the fiber/base distinction sharp — SS distinguishes the two motions, Theorem 6.2 twins PhC (fiber) with ALP (base), and one cannot twin two things that are not two. But 6.3B *dissolves* that distinction. Does the conjecture saw off the branch the theorems sit on? The proposed reconciliation is a stratification by scale: the distinction *resolved (sharp) above the distinguishability floor*, where the theorems live and the twinning is statable, and *dissolved at the floor*, where the transfer happens and the lift crosses what was, above the floor, a sector boundary — one space at two scales, π looking fibered above the floor and unfibered at it. We previously presented this as a *mechanism*. The referee is right that, stated as bare scale-talk, it is engineered to the result — sharp exactly where we need it sharp, unreal exactly where we need it unreal, reconciled by stipulation. So we downgrade it, and offer a toy model to earn at least its *coherence*.

Toy model (coherence of a scale-dependent involutive→bracket-generating distribution — illustration, not actuality). Take $\mathcal{A}\_toy = B \times T^2$ with $B = [0,1]$ (carving-parameter, the base) and T^2 the two-dimensional phase-fiber, coordinates (x; θ, φ), and π the projection to x. A rank-2 distribution needs a *third* dimension for brackets to point somewhere new — on a 2-manifold a rank-2 distribution is all of the tangent space trivially and a single base-transverse field can never bracket-generate, so the base-transverse field must couple to a phase direction *other* than the one already spanned. Equip $\mathcal{A}\_toy$ with

$$D = \mathrm{span}\{ \partial\_\theta , \partial\_x + f\_\varepsilon(x,\theta)\, \partial\_\varphi \},$$

coupling the base direction to φ (not back into θ, which is already in D). Then

$$[\, \partial_\theta\, ,\, \partial_x + f_\varepsilon\, \partial_\varphi \,] = (\partial_\theta f_\varepsilon)\, \partial_\varphi.$$

Gate the coupling by scale: f_ε *constant in θ above resolution* (so $\partial_\theta f_\varepsilon = 0$, the bracket vanishes, D is **involutive** and integrates by Frobenius to a foliation — the phase-fibers are sealed, no admissible direction is base-transverse, the ALP-analogue fails) and *θ-dependent at the ε-grain* (so $\partial_\theta f_\varepsilon \neq 0$, the bracket $(\partial_\theta f_\varepsilon)\partial_\varphi$ is transverse to D, and D + [D,D] spans $T\mathcal{A}_{toy}$ — D is **bracket-generating**, Chow–Rashevskii connects across fibers, the ALP-analogue holds). Above the floor D is involutive (fibered, ALP-analogue fails); at the floor D bracket-generates (connected). This is the textbook Martinet/Heisenberg construction — a rank-2 distribution in a 3-manifold whose bracket switches on — and the involutive↔bracket-generating switch is exactly the stratification's content, with no ill-defined tolerance-quotient and no collapse.

This toy is well-defined and survives contact: a smooth distribution with a scale-gated coupling to a transverse direction, the standard sub-Riemannian setting, not a quotient by a within-ε tolerance relation (which, being non-transitive, either fails to be a quotient or collapses the space to a point under ε-chains — the wrong vehicle). It proves the stratification's **coherence, not its actuality**: we do not claim $\mathcal{A}$'s admissible distribution *is* this D, only that "a distribution involutive above scale and bracket-generating at the floor" is a consistent structure exhibiting precisely the scale-dependent transfer the stratification needs.

One structural caveat the construction makes explicit: the mechanism **requires fiber dimension ≥ 2** — a transverse phase direction (φ) for the bracket to land in, distinct from the one already spanned. On a one-dimensional phase fiber (a single S^1) the bracket would have nowhere transverse to go, the 2-manifold trap recurs, and D could only be involutive or trivially all of $T\mathcal{A}$. The running gloss elsewhere — "the phase S^1," PhC as a one-parameter torus action, "a sub-resolution phase displacement" — reads as effectively one-dimensional phase; but §2.4 gives the fiber as "the amplitude/phase data H0 assigns at that carving," which need not be one-dimensional. Whether $\mathcal{A}$'s actual fiber supplies a phase dimension ≥ 2, rather than a single phase circle, is a structural question the construction must eventually settle: the toy's mechanism presupposes it. (The toy proves coherence regardless — it exhibits *a* space with the switch — but its mechanism transfers to the actual $\mathcal{A}$ only if the actual fiber has the transverse room, which we flag rather than assume.)

We label OSP **proven modulo the programme's operationalism**, 6.3A-bridge and 6.3B **conjectural**, and the stratification **conjectural-with-toy-coherence**. The previous paper's discipline forbids treating "π is single-topology" as established because it is plausible and the default now favours it. §7 says why, even with the burden shifted, we cannot assume it.

---

7. The non-circularity test: when General Continuity reduces RC, and when it renames it

Call the single-topology hypothesis **General Continuity (GC)**: *the admissible-direction distribution D is bracket-generating, so admissible motion connects across fibers (Chow–Rashevskii) and base paths admit admissible lifts.* The seductive form of Thesis B is: GC ⟹ ALP, ALP ⟹ RC ($\mathscr{R}$-connectedness entailed, §3), so GC ⟹ RC, arc closes. The question is whether GC is a genuine *reduction* of RC or RC *renamed*.

The test is the one *No Contextual Weight* used to separate "w is a measure" (a genuine reading of N4) from "subdivision-invariance" (which, given separability, simply *was* PC-C):

Non-circularity test. GC reduces RC only if GC can be *established by an argument that does not route through RC itself* — only if there is an independent reason for π's single-topology that does not assume base paths admit admissible lifts. If the only route to GC is to assert that base paths are admissibly lifted, then GC *contains* ALP, RC's dynamical residue, and "GC ⟹ RC" is "RC ⟹ RC."

Applying it:

- **The $\mathscr{R}$-connectedness half passes the test (and is cheap).** §6.1 establishes path-connectedness of the base from the transport topology + capacity-continuity — facts about the *base's* geometry, established independently of any claim that motion lifts paths. A connected base is compatible with a base whose paths admit *no* admissible lift (the fibered horn). So $\mathscr{R}$-connectedness is non-circular — and, as §6.1 stresses, cheap; the decomposition's value is isolating it so the difficulty concentrates wholly in ALP.
- **The ALP half is where circularity threatens, and where the test must be passed or the gate stays shut.** To get ALP from GC non-circularly, π's single-topology must follow from something that is *not* "base paths are admissibly lifted." The OSP ∧ 6.3A-bridge ∧ 6.3B route attempts exactly this: it derives single-topology from operational individuation (OSP) and a distinguishability fact about *configurations* (6.3B), not from any premise that lifts exist — so *if it goes through*, it passes the test. Until it does, we have no non-circular route to GC's lift content, and must not help ourselves to GC. Asserting "π is single-topology, lifts throughout" without the OSP/distinguishability argument is precisely asserting ALP, RC's residue — the circle.
- **The split introduces circularity surfaces, and we guard them — at two depths plus the OSP scope.** Because the route runs through 6.3B and OSP, the test bears on their *form*. *First depth (executed-motion reading of 6.3B).* 6.3B passes only as indistinguishability of motion-*directions/configurations*; read as "refinement motion *is executed* indistinguishably from phase motion," it presupposes base lifts — a sliver of ALP — and is circular. *Second depth (existence-of-refinement-motion, subtler — and sharper in the distribution language).* "Sub-resolution re-carving" can presuppose the

substrate *can* re-carve — that refinement *motion* exists to have a sub-resolution instance. If admissible refinement *motion* (not just refinement *configurations*, which §6.1 supplies) is built into 6.3B's setup, the conjecture assumes a form of ALP en route to delivering it. The guard, in §6.1's vocabulary: 6.3B compares two *displaced configurations* and asserts their indistinguishability (6.3B-config); it must not presuppose a *motion* producing the displaced carving, and its proof must *derive* sub-resolution refinement motion from distinguishability and H0. The distribution framing makes this guard sharper and harder: "D is bracket-generating" (6.3B-dist) speaks of admissible base-transverse *tangent directions*, so the translation 6.3B-config ⟹ 6.3B-dist is *itself* the place the existence of admissible refinement directions would be smuggled if one is careless (§6.3). The guard therefore extends: the eventual proof must derive 6.3B-dist *from* 6.3B-config — establish that operational indistinguishability of displaced configurations *yields* admissible base-transverse tangent directions — without assuming those directions exist. Asserting "D bracket-generates" directly, rather than deriving it from configuration-indistinguishability, is the second-depth circularity in distribution clothing. *Third surface (OSP over-collapse).* OSP must be held to sector boundaries, not intra-class distinctions (the §6.3 selectivity guard); an unrestricted OSP would dissolve distinctions the programme keeps, and a critic recategorising the fiber/base boundary as intra-class would block the route — that recategorisation, not unrestricted verificationism, is the live objection. Held to all three, single-topology is derived from the *structure* of $\mathcal{A}$ and the operational individuation the programme already runs on, not from any premise that motion lifts paths — and the circle stays open.

So the discipline yields a clean verdict on what closure requires: **Thesis B holds iff (i) PhC is proven (§4 tier), and (ii) π's single-topology is established by an argument independent of RC — for which OSP ∧ 6.3A-bridge ∧ 6.3B is the candidate, with 6.3B held to its directions-only form (first depth) and proven to derive rather than presuppose refinement motion (second depth), and OSP held to sector boundaries (third surface).** Neither is in hand. What *is* in hand: RC is not new in kind (Thesis A); its base half is discharged cheaply and non-circularly ($\mathscr{R}$-connectedness); its dynamical half is PhC's twin (ALP) with the transfer isolated to one lift-property of π, the non-circular route identified-but-not-walked at all three surfaces, and the burden on the fibered horn shifted (§6.3) — without the gate being declared shut.

Remark 7.1 (we have not closed the arc, and say so). It would be easy, and wrong, to write the previous sections up as "RC reduces to phase continuity, arc closed." That is the symmetric fiat the programme guards against — proving structural facts (SS, the decomposition, $\mathscr{R}$-connectedness, the twinning) and then *announcing* the dynamical conclusion. We prove the structural facts and stop at the gate. Thesis B is conditional on the phase tier and conjectural on single-topology; we leave it there. And note (§10.5) that "not new in kind" is compatible with "just as unprovable" — localization is not reduction, and we do not let the rhetoric outrun that.

---

8. Why this is not forcing the gate

The paper makes claims of three strengths, and their separation is the discipline.

What is proven (no fiat). Same-Species (Theorem 5.3): phase and refinement motion share all admissibility credentials and differ in none *of them* — a classification on the data axis, settled as IA was (they differ in fiber-vs-base locus, §5.6, a difference within the class, not of membership). Thesis A (Corollary 5.4): RC is therefore the same *kind* of condition as PhC. The RC factoring (Proposition 3.1): RC ⇔ ALP, with $\mathscr{R}$-connectedness entailed by ALP (a liftable path is a path) and factored out only to show the difficulty does not lie in it — not a splitting into independent halves. ALP ≡ PhC-for-refinement (Theorem 6.2): the dynamical content is PhC's twin — same predicate, two path-families — by SS. OSP (Proposition 6.3.0): sector boundaries need an operational witness, modulo the programme's operational individuation. None of these asserts that any base path is in fact admissibly lifted; none closes the arc.

What is cheap, not a theorem (honestly demoted). $\mathscr{R}$-connectedness (Proposition 6.1, via Lemma 6.0): under the transport topology of §2.4, path-connectedness of the base is near-automatic from capacity-continuity, *modulo* the asserted (not proved) judgement that Vol_op carries no internal refinement-separating invariant. We no longer present this as a substantial result — the transport metric makes it near-definitional (Remark 6.0.1) — and the demotion is the point: it concentrates the entire difficulty in ALP, where it belongs. The metric is a flagged choice (§2.4), adopted as the minimal topology generated by Vol_op, precisely so connectedness is not engineered into the open sets.

What is conjectural, and explicitly not assumed. π's single-topology — that base paths admit admissible reversible lifts (ALP) — derived from OSP ∧ 6.3A-bridge ∧ 6.3B. OSP is near-definitional (operational individuation, §6.3); 6.3A-bridge (the motions are the complete witness-set for the fiber/base boundary) and 6.3B (the motion-directions are indistinguishable at the floor) are the conjectural content, and now the precise locus of debate. The split gives a referee sharp targets — reject OSP's application to this boundary, exhibit an independent witness against 6.3A-bridge, or deny 6.3B's indistinguishability. The burden-shift (§6.3), an application of OSP + 6.3A-core, makes the fibered horn owe an *independent operational witness* of the boundary. We label OSP proven-modulo-operationalism, 6.3A-bridge and 6.3B conjectural, the stratification conjectural-with-toy-coherence (Remark 6.3.1); hold 6.3B to its directions-only form and OSP to sector boundaries so the route stays non-circular (§7); and walk no further toward declaring the gate shut.

The asymmetry to guard, inherited and now turned on ourselves: the previous papers warned against a *negative* result leaking into a *positive* dynamical conclusion, and against an *intra*-refinement fact leaking into a *cross*-refinement claim. Here both temptations point the same way — toward declaring closure — because we *want* PhC (fiber) to deliver ALP (base lift). The fiber–base gap (§6.3) is that very leak, and §7's refusal to cross it without the OSP/distinguishability argument is the discipline holding. We open RC — decompose it, demote the base half, twin the lift residue with PhC — but we do not let the twinning *announce* the lift. The gate stays exactly as open as the argument leaves it.

---

9. Assembly: RC localized, not lessened

Across the now-six companions the Born exponent stands at:

$\ell^2 \Leftarrow$ IA $\wedge$ Obstruction B, IA proven, Obstruction B affirmative $\Leftrightarrow$ RC, RC open.

This paper opens RC, in the constructed spaces of §2.4:

1. **RC ⇔ ALP, $\mathscr{R}$-connectedness entailed** (§3, proven) — RC is the Admissible Lift Property; $\mathscr{R}$-connectedness is entailed by ALP and factored out only to show the difficulty is not in it.
2. **SS: phase and refinement motion are one class** (§5, proven modulo inherited criteria) — whence **Thesis A**: RC is not a new *kind* of dynamical debt; it is the refinement-axis instance of the phase-axis continuity.
3. **$\mathscr{R}$-connectedness — cheap under the transport topology** (§6.1, Proposition 6.1 via Lemma 6.0; near-definitional, modulo the asserted no-internal-invariant judgement) — *not* a substantial theorem; the demotion concentrates the difficulty in ALP.
4. **ALP ≡ PhC-for-refinement** (§6.2, proven via SS) — the lifting residue is phase continuity's twin, the same lift-predicate on base paths instead of fiber paths.
5. **The fiber–base gap, as a lift property of π** (§6.3, §7) — PhC delivers ALP iff π is single-topology (base paths admissibly lift); derived from OSP ∧ 6.3A-bridge ∧ 6.3B, burden shifted onto the fibered horn, the whole refused as an assumption.

The picture that replaces "RC, one monolithic open dynamical gate":

- **Base content of RC:** $\mathscr{R}$-connectedness — cheap under the transport topology, near-automatic, *not* the locus of difficulty.
- **Lifting content of RC:** ALP — the twin of PhC. **Reduced in kind** from "new continuity premise" to "the refinement-sector instance of the lift property the programme already owes," but **not reduced in difficulty**: whether π admits the lift is exactly as open as RC was.

So the honest headline is **RC localized, not lessened.** SS, the decomposition, the demotion, and the twinning relocate the difficulty and classify it; they do not diminish it. The entire difficulty now sits, undiminished, in one precisely-stated place: does π admit an admissible reversible lift of base paths (ALP) — that is, is the admissible-direction distribution bracket-generating, and, if "reversible" demands it, bracket-generating *with a measure-preserving lift*, the strictly-stronger second tier §10.5 leaves open and the burden-shift does not reach? That is a real gain — a monolithic, unexamined dynamical gate has become a precisely-tiered lift-property of a constructed projection, shown to be the same species as a property the programme already faces on the phase axis — but it is localization and classification, not reduction. The residue on which closure waits is the conjunction of *PhC proven* (§4 tier) and *the lift-property* (bracket-generation, possibly with measure-preservation; §6.3 and §10.5 conjectures), the latter carrying the whole weight.

Under the optimistic reading (PhC proven; single-topology established, via OSP ∧ 6.3A-bridge ∧ 6.3B): GC holds, GC ⟹ PhC ∧ ALP, and with $\mathscr{R}$-connectedness and IA,

$\ell^2 \Leftarrow$ IA ∧ GC, IA proven, GC the single lift-property of admissible motion.

and the squaring closes. What had looked like a premise beyond FP1's reach — RC, a dynamical fact the founding rule is silent on — is then supplied by one lift-property, owed once and holding on both fiber and base paths: FP1's silence on dynamics is supplemented not by a new posit but by the single continuity the programme already leans on.

Under the guarded reading (PhC posited, or single-topology open): the arc does not close, and RC's difficulty is undiminished — but it is no longer a mysterious monolith. It is one lift-property of π, of the same kind as the phase-axis property the programme already runs on, with the base half demoted to cheap and the lifting half twinned and precisely stated — and the open work sharply specified: prove PhC, and establish that π admits admissible base-lifts without assuming it.

Remark 9.1 (what changed, precisely). The previous assembly carried RC as the sole open gate, flagged "plausibly not new, the refinement-axis kin of the phase continuity already owed." This paper converts the parenthetical into structure: SS proves the *kinship* (same class); the construction of π, the decomposition, and the transport-topology demotion show the base half is *cheap, not the difficulty*; Theorem 6.2 shows the lifting half is PhC's *twin*; and §6.3–7 isolate the *one* question — does π admit admissible base-lifts (single-topology) — that the kinship does not settle, refusing to settle it by fiat. RC is not retired (we do not prove it) and not lessened (the lift question is as hard as RC was); it is *localized* — from "an unexamined dynamical gate" to "one named lift-property of a constructed projection, the same species as a phase-axis property already owed." The squaring waits on whether that one lift-property holds.

---

10. Limitations and Open Problems

10.1 The phase tier — proven or posited (decides Thesis B's reachability)

PhC (§4) is leaned on by the Diagonal-Torus selection. Whether its *continuity* is derived within the selection or built into the choice of a torus decides whether closure is even reachable on the phase side. **Conditional**, to be confirmed against the linearity companion's Diagonal-Torus construction. If posited, this paper delivers Thesis A only; if derived, Thesis B is reachable subject to §10.2. This is the first thing to check, and Proposition 4.1 shows it cannot be finessed: any direct proof of ALP *is* a proof of PhC-for-refinement and settles the tier rather than avoiding it.

10.2 Bracket-generation of the admissible-direction distribution (π "single-topology") — the open gate (decides Thesis B, given 10.1)

Whether the admissible-direction distribution D is involutive-and-fiber-tangent (PhC does not transfer; ALP fails) or bracket-generating (base paths admissibly lift; ALP holds) is the one genuinely open structural question — now a precise property of D, located in the admissible directions, not in the point-set topology of $\mathcal{A}$ (§6.3, §10.5). We use "π single-topology" only as informal shorthand for "D bracket-generating." We derive bracket-generation from OSP ∧ 6.3A-bridge ∧ 6.3B. OSP (Proposition 6.3.0) is **proven modulo the programme's operational individuation**, held to sector boundaries by the §6.3 selectivity guard. 6.3A-bridge (the motion-directions are the complete witness-set for a fiber-tangent involutive segregation of D) and 6.3B (the fiber and base directions are operationally indistinguishable at the floor, so D is bracket-generating there) are **conjectural** and are now the precise locus of debate — a sceptic must accept OSP and 6.3B yet *exhibit an operational witness, independent of the motions, that D is fiber-tangent involutive*, or deny 6.3B. The natural next technical step is to prove 6.3B from the distinguishability floor and H0 — held to its directions-only form (§7) — and to establish 6.3A-bridge's completeness claim. Success closes the gap (given 10.1); failure — a positive exhibition of an independent involutivity-witness, i.e. that D genuinely foliates into the fibers — would establish ALP as an independent dynamical premise and locate RC's irreducible dynamical core precisely. Either outcome is a result; the burden-shift (§6.3) makes the failure case the one owing the positive construction.

10.3 $\mathscr{R}$-connectedness is cheap, and its one asserted premise

$\mathscr{R}$-connectedness (§6.1, Proposition 6.1 via Lemma 6.0) is **not a substantial theorem**: under the transport topology of §2.4 it is near-definitional (Remark 6.0.1). Its one non-trivial dependence is the **asserted, not proved** judgement that Vol_op carries no internal refinement-separating invariant (the external counterpart already excluded by the prior papers). We judge it satisfied — capacity is a single additive content (CA), offering no obvious conserved structure to separate refinements — but this is a plausibility note, not a demonstration that CA admits no such structure; a sceptic of $\mathscr{R}$-connectedness must *exhibit* an invariant, and we would have to *rule one out*, which we have not done rigorously. Capacity-continuity is the other dependence, judged inherited. The honest status: connectedness is cheap and we lean on it lightly, with the no-internal-invariant judgement flagged as the one place it could fail.

10.4 Whether single-topology and the phase tier are one question

6.3B, if proven, would establish single-topology *via* the substrate's operational individuation at the floor — and that same individuation is the natural candidate for *deriving* PhC (§10.1) rather than positing it. So §10.1 and §10.2 may be one question: does finite distinguishability force the admissible-direction distribution to be bracket-generating (π "single-topology"), with both phase-continuity and refinement-lift as instances? If it does, it discharges both the phase tier (PhC derived) and the gap (single-topology), and Thesis B closes in one stroke. This reconciles §4's guarded fork with this hope: §4 treats "PhC posited" as live and perhaps actual (Thesis A only), while this unification would *retroactively settle §4's fork in favour of "derived"* — the two in sequence, not tension. **Conjectural** that they coincide; **plausible**, both being the lift-property of admissible motion at the floor. The cleanest closure and the natural unifying target.

10.5 The connection/holonomy sharpening (the strongest available form)

The cleanest form of the open gate, which we flag as the next technical target rather than assert, is sub-Riemannian. Because admissible motion is a *restricted curve class* (§2.4), the natural object is the distribution $D \subset T\mathcal{A}$ of admissible directions (§6.3), and ALP — base paths admit admissible lifts — is a **controllability** question for D, governed by **Chow–Rashevskii**: admissible motion connects across fibers iff D is **bracket-generating**. The fibered horn is exactly D *involutive and fiber-tangent* — a foliation by phase-fibers (Frobenius), no admissible direction with any base component. So the obstruction is the **involutivity-vs-bracket-generation** dichotomy of D, a named object in sub-Riemannian geometry, and 6.3B's content becomes the precise claim *finite distinguishability makes D bracket-generating at the floor*. This is sharper than the bundle-section or point-set-topology framings and gives 6.3B a theorem-class (sub-Riemannian controllability) to aim at.

A correction to the natural first guess (holonomy is the wrong invariant). One is tempted to say: if admissible motion is *horizontal lift* with respect to a connection on π, then ALP is the horizontal-lift property and the obstruction is the connection's holonomy. This is wrong, and the way it is wrong is instructive. For a genuine Ehresmann connection, horizontal lifts of base paths *always exist* — that is what a connection is *for*; holonomy obstructs whether a lift around a *loop* closes up, not whether lifts exist. So if admissible motion were horizontal lift, ALP (which only asks for *some* liftable path) would be free, and the gate would evaporate — "too good," which is the tell that holonomy mislocates the obstruction. The obstruction lives one level down: whether the admissible class is rich enough to *be* a connection — to have any base-transverse component — versus being foliated into the fibers. That is precisely the bracket-generating/involutivity dichotomy, not holonomy. We therefore frame the sharpening around Chow–Rashevskii and D's involutivity, not around a connection's holonomy.

We do **not** assert the differential apparatus: we have not shown the substrate's admissible-motion structure supplies a smooth distribution D with well-defined brackets (it may not, and we will not invent differential geometry the programme cannot cash — Chow–Rashevskii presupposes a smooth manifold and a smooth distribution). We flag it **conjectural** — *if $\mathcal{A}$ carries enough structure for an admissible-direction distribution D, then ALP is D's controllability and the gate is D's involutivity vs bracket-generation, with 6.3B the claim that finite distinguishability makes D bracket-generating at the floor* — as the strongest and most testable reformulation, and the natural next step after 10.2.

A sequel to the correction (Chow–Rashevskii gives connectivity, not measure-preservation — so the gate may be strictly stronger than bracket-generation). The holonomy correction has a sequel that tightens the open problem, and it stems from our own admissibility criterion. Chow–Rashevskii delivers *connectivity* by admissible paths, but the sub-Riemannian lift can have length (and measure-distortion) far exceeding the transport distance — bracket-generation buys a path, not an isometry. Now recall that admissible motion is required to be **capacity-preserving and reversible** (§2.4). "Reversible" plausibly demands more than bare connectivity: a measure-preserving, or at least controlled-distortion invertible, lift — not merely *some* horizontal path. If so, ALP is not "D bracket-generating" (controllability simpliciter) but "**D bracket-generating with an admissible measure-preserving lift**," controllability *within the measure-preserving sub-class*, which is strictly stronger and which Chow–Rashevskii does not deliver. This sharpens the open gate: bracket-generation is then *necessary but possibly not sufficient* for ALP. The same point meets §10.6 from the companion side: if *No Pre-Individuation*'s RC carries a uniformity or measure-preservation clause beyond bare traversal, bracket-generation alone cannot supply it, because the sub-Riemannian lift may distort measure wildly. Either way — from our own "reversible" criterion or from the companion's RC — the honest statement of the gate is bracket-generation *plus* an admissible measure-preserving lift, and we flag the measure-preservation component as a distinct, strictly-stronger conjectural requirement that the next technical step must address, not fold silently into controllability. We had understated this by equating ALP with controllability; corrected, ALP inherits whatever measure-preservation "reversible" demands.

This strengthening interacts with the §6.3 burden-shift, and we surface the interaction rather than leave it unremarked. The burden-shift (§6.3) voids the *directional* sector boundary — base-transverse vs fiber-tangent admissible directions — by showing it operationally witnessless (OSP + 6.3A-bridge + 6.3B). If the operative gate is now the *measure-preserving* one, the relevant boundary shifts to "directions admitting measure-preserving lifts vs directions admitting only distorting lifts," and one might worry that measure-distortion — which the programme's capacity-bookkeeping *does* register, as a change in apportioned content along the path — supplies exactly the independent operational witness the directional boundary lacked, handing the fibered horn what §6.3 claims it cannot find. Two responses, and we give both at their proper tiers.

First (which we judge correct, on §6.3's own selectivity guard): capacity-distortion does **not** supply a *sector* witness, because it is *intra-class*. Distortion is a property comparing one admissible motion to another by how much content each shifts en route — a distinction *among* admissible motions, witnessed by the motions differing, not a partition of the admissible class into admissible and inadmissible directions. OSP's selectivity guard explicitly sets aside exactly such intra-class distinctions (§6.3); so capacity-bookkeeping registering distortion is no more a sector witness than two admissible configurations differing was. The burden-shift therefore *survives for what it was for*: voiding the directional boundary.

Second (the honest limit): the burden-shift establishes only the **controllability** gate (D bracket-generating), and its extension to the strictly-stronger **measure-preserving** gate is **open**. This is not because the strengthening hands over a witness — the first response forecloses that — but because the burden-shift was never aimed at the measure-preserving gate: bracket-generation voided of directional obstruction gives *a* lift, not a measure-preserving one. Whether a measure-preserving lift then exists is a separate question the burden-shift does not reach. So: the burden-shift is established for the controllability gate; the measure-preserving gate is a further open requirement, neither delivered by the burden-shift nor (by the intra-class argument) obstructed by an independent witness the strengthening supplies. We flag it as open and distinct, next to the sequel that raised it.

10.6 The exact form of RC against *No Pre-Individuation*

We reconstructed RC as the affirmative resolution of Obstruction B and stated it as $\mathscr{R}$-connectedness ∧ ALP (§3). The reconstruction should be checked against *No Pre-Individuation* §7.3's exact statement: if RC there carries content not captured by $\mathscr{R}$-connectedness ∧ ALP — for instance a uniformity or measure-preservation clause on the lifting motion beyond bare admissible lift — that clause must be added and its status separately determined. We judge the two-part decomposition faithful, but the companion is authoritative.

10.7 "Not new in kind" is compatible with "just as unprovable" — localization is not reduction

We state plainly the limit of Thesis A. That RC is the same *kind* of condition as PhC (proven) is compatible with RC being exactly as unprovable as it was before — unprovable-and-of-the-same-species is still unprovable. The honest gain of this paper is **localization and classification**, not reduction of difficulty: the difficulty is concentrated, undiminished, in the single lift-property of π (§9). We have tried throughout not to let the rhetoric ("the last gate," "mostly already turned") outrun this; where earlier drafts said "RC shrinks," the accurate claim is "RC is localized." A reader should take from the paper that the open question is now sharp and singular, not that it is smaller.

10.8 Internal coherence is not empirical adequacy (the gap a journal referee will care about most)

This paper is an internal-consistency reconstruction: it determines RC's status *within* the programme's framework. It says nothing about whether the framework tracks physics. Given that the surrounding programme claims to derive the Born rule, the Standard Model gauge group, and α^{-1} to 15 ppm, the gap between "internally coherent reconstruction" and "correct about the world" is the one a journal referee will weigh most heavily — and it is one this paper neither closes nor claims to. Localizing RC to a lift-property of a constructed π is progress in the programme's internal logic; it is not evidence that the construction describes nature. We record this as a scope limitation, not a defect, because the paper's question is explicitly internal.

10.9 The constructed spaces are this paper's, adopted minimally

$\mathscr{R}$'s transport metric, $\mathcal{A}$, π, and the admissible-motion class (§2.4) are not inherited from prior work; we adopt them as the minimal structure generated by Vol_op and flag every choice (§2.4 disclosure). A programme-canonical topology or specification of π, fixed later, could redistribute the load between connectedness and lifting — though the transport choice was made precisely to push difficulty onto lifting rather than hide it in the topology. The construction should eventually be reconciled with the *Packing* papers' treatment of Vol_op and admissible refinements.

10.10 No data-axis content is reopened

Nothing here reopens the data axis the previous paper closed. SS is a classification (data-axis, like IA); $\mathscr{R}$-connectedness is base-geometry (data-axis). The genuinely dynamical residue (ALP) is kept strictly separate and is not claimed proven. We move the base half of RC onto the data axis by argument and leave the lifting half on the dynamics axis, localized but open.

10.11 The Admissible Lift Property as the next theorem, stated on its own

The paper has, throughout, been circling a single question, and it is worth stating it as a *standalone problem* — the programme's next named theorem — rather than only as RC's residue. Decoupled from RC, it reads:

The Admissible Lift Problem. Given the projection $\pi : \mathcal{A} \to \mathscr{R}$ (the motion space over the carving space, §2.4) and the restricted class of admissible reversible motions, *when does a continuous path in $\mathscr{R}$ admit an admissible reversible lift?* Equivalently, in the distribution language: when is the admissible-direction distribution $D \subset T\mathcal{A}$ bracket-generating across the fibers — and, if "reversible" demands it, bracket-generating with a measure-preserving lift?

This is no longer about RC's status; it is a clean question in sub-Riemannian / controllability geometry, posed for the substrate's admissible-motion structure. Everything in this paper is, in effect, the claim that *this* is the right question — that RC, the Born exponent's last dynamical conjunct, reduces to it (RC ⟺ ALP, §3); that it is the same species as the phase-axis continuity already owed (Theorem 6.2); that its obstruction lives in D's involutivity-vs-bracket-generation (§6.3, §10.5), not in point-set topology; and that finite distinguishability is the candidate engine for its affirmative answer (§6.3, conjecturally). We flag the Admissible Lift Problem as the natural successor target: a problem stated in standard geometric terms, with a named theorem-class to draw on, whose resolution (with the measure-preserving tier, and given the phase tier of §4) would close the squaring — and whose negative resolution would locate RC's irreducible dynamical core precisely. The contribution of *this* paper is to have turned "is RC a new premise?" into "does π admit admissible lifts?" — and that reformulation, more than any single result here, is what we take forward.

---

11. Conclusion

The squared-amplitude rule, chased through five companions, came down to one open gate: RC, whether the carvings flow — a dynamical continuity premise the founding rule cannot reach, flagged as plausibly the refinement-axis kin of a phase-axis continuity already owed. This paper opened it, with the goal not of proving the Born rule but of determining whether RC is genuinely new.

It is not new in *kind*. Phase motion and refinement motion share every credential the substrate uses to admit a motion — admissible, reversible, pre-commitment — and differ in none of them; they are one class. (They differ in one respect, fiber-vs-base locus, which is what §6 turns on — but that is a difference within the class, not a difference of admissibility.) So RC is the refinement-axis instance of the phase-axis continuity, not a second species of debt. That much is proven, on the data axis, the way internality was.

And RC is not even monolithically dynamical. Opened up in the constructed spaces, it splits along the line between a configuration and a motion: that the carvings form one connected landscape — that a continuous family of admissible *carvings* interpolates any two — is, under the natural transport topology, a near-automatic fact, with no motion in it; the base half of RC, and cheap. What remains is the *lift* of that landscape by admissible *motion*, and that is the exact twin of the phase continuity: the same predicate — admits-an-admissible-reversible-lift — asked of the paths between carvings instead of the paths within one. The two do not overlap: a connected landscape of admissible configurations can still be one no admissible motion lifts, which is exactly the gap that remains.

So the last gate is a familiar *kind* of lock — but we refuse the step that would call it turned. Phase continuity holds *within* a carving; the lift is *between* carvings; and the previous paper's hardest-won lesson was that a within-carving fact does not reach a between-carving claim for free — the very leak, now pointed at the conclusion we want. Whether the phase continuity carries across that gap depends on a single structural fact about the projection π: whether it is single-topology (base paths admissibly lift) or fibered (they do not). There is a substrate reason to expect the former — finite distinguishability cannot tell a sub-resolution re-carving-direction from a sub-resolution phase-direction, so by operational individuation (OSP) it should forbid the two from sealing into separate sectors — but that reason is a conjecture (6.3A-bridge and 6.3B), not a proof, and to assume it would be to assume the lift, which is to assume RC. We do not.

RC, then, is **localized, not lessened.** Not retired — we have not proven the carvings flow — and not made smaller: the lift question is exactly as hard as RC was. What is gained is precision and singularity. RC is not new in *kind* (same admissibility class as PhC, proven); its base half is cheap (near-automatic under the transport topology); and its dynamical half is one named lift-property of a constructed projection, the same species as a phase-axis property the programme already owes. What stands between finite distinguishability and the Born rule is no longer a monolithic dynamical mystery but two named things: whether the phase continuity is proven rather than posited, and whether π admits admissible lifts of base paths — bracket-generation of the admissible-direction distribution, and, if "reversible" demands it, the strictly-stronger measure-preserving lift §10.5 leaves open. If these hold — and the phase tier and the lift may be one question, answered together by finite distinguishability at the floor — then there is one lift-property, owed once, and the squaring closes. Whether this internal closure would describe nature is a further question, and not one this paper can reach.

Whether the carvings flow now waits on one thing of its own kind: whether the projection is single-topology, or two.