The Admissible Lift Theorem

Bracket Generation, Sealed Leaves, and the Final Reversible-Connectedness Test

Keith Taylor — VERSF Theoretical Physics Programme

General Reader Summary

A long line of work in this programme has traced the origin of quantum probability back to a single yes-or-no question. Before reality "commits" to a definite fact, it can be described in many partial, not-yet-settled ways — call these *refinements*. The question is whether these partial descriptions can all be continuously reshaped into one another, or whether some of them are walled off from the rest.

If they can all reach one another, the framework produces the behaviour quantum mechanics needs. If some are walled off, reality splits into separate compartments and a more classical picture survives instead. Everything in the quantum-reconstruction effort now hinges on which of these is true.

Earlier papers circled this question from several angles without ever attacking it head-on — a real danger, because a research programme can look busy while only renaming its central problem. This paper attacks it directly, and the strategy is to keep asking: *if* the partial descriptions are walled off, what exactly could be doing the walling? Any such wall would have to be some hidden quantity that stays fixed no matter how you reshape a description — a "seal." The bulk of the paper is the systematic elimination of every kind of seal that could exist.

One by one, the candidates fall. A seal can't be an external label (the framework forbids labelling descriptions from outside). It can't be built from the sizes of the parts (those can always be continuously adjusted). It can't be the ordering of the parts, or the number of parts, or any other feature of *how* a description happens to be drawn — because all of those can be changed by legitimate reshaping. What's left, after everything local and superficial is stripped away, is only the *global shape* of the space of descriptions: whether it comes in separate pieces, or has "loops" in it that can't be undone.

So the paper proves the hard half of the question: there is no built-in wall. Every ordinary candidate for a seal is ruled out. The only way reality could still be walled off is if some hidden, global, shape-based quantity were *added on top* of the founding assumptions — a quantity that, by construction, would do no observable work. Whether such a thing exists is the one question left genuinely open; the paper argues it shouldn't be assumed without reason, while being honest that this is a judgement of economy, not a proof.

The payoff is twofold. First, the central question is no longer "why should reality be connected?" but the far narrower "does this one specific, well-defined structure carry a hidden global feature?" — a question that can in principle be computed rather than argued. Second, that surviving global feature turns out to be the *same kind of object* studied in a separate strand of the programme (the Gate-3 closure work), so two problems that looked independent are revealed to be, at minimum, the same kind of problem — possibly the same problem. The paper builds that bridge carefully, by elimination, without assuming it.

*Epistemic markers: **(established)** for results inherited from prior VERSF papers or standard mathematics; **(assumption)** for conditions adopted explicitly; **(conditional)** for results that hold under a stated assumption; **(open)** for what remains undecided.*

Table of Contents

Abstract

Previous papers reduced the Born-rule residue to Reversible Connectedness (RC), then reformulated RC through the Admissible Lift Property (ALP), operational sector analysis, admissible motion spaces, and refinement topology. The present paper argues these reformulations should not be treated as independent open problems: they are consequences of a deeper *local* question about the generator structure of admissible reversible refinement motion.

We introduce the admissible-direction distribution D on refinement space $\mathscr{R}$ and the Lie algebra Lie(D) generated by admissible infinitesimal motions. A fork theorem is established: under a stated regularity assumption, if D is bracket-generating on a connected $\mathscr{R}$ then the Chow–Rashevskii theorem yields ALP and hence RC; if Lie(D) closes into a proper involutive sub-distribution of constant rank, then $\mathscr{R}$ foliates into dynamically sealed leaves and RC fails. The singular (non-constant-rank) case is handled separately via the Stefan–Sussmann orbit theorem, which replaces the clean foliation with a partition into orbits of possibly varying dimension.

The entire RC problem is thereby reduced to a generator question, and two further results sharpen it. First, the admissible generators are realised concretely as capacity transfers on a simplex; their Lie brackets vanish, so bracket generation collapses to connectivity of the admissible-transfer graph. Second, a No Internal Seal Lemma exhausts what could disconnect that graph: an external label (excluded by IA), a capacity-derived invariant (proved mathematically constant, hence no seal), or a dynamical conserved charge (which cannot come from FP1 and must be imported as operationally idle extra structure). The negative half of RC is thereby proven — there is no FP1/IA-native obstruction — leaving a single residual question: does the admissible-motion dynamics conserve a sector charge beyond FP1? The paper does not settle that question; its contribution is to corner RC's only possible obstruction into one non-native, operationally idle dynamical charge, with the parsimony case for its absence stated honestly as defeasible. Under full merge/split equivalence the result sharpens further. RC has two senses that must be separated: reachability (RC_reach, an H^0 property) and path-independence (RC_path, an H^1 property). The component obstruction is eliminated — Q_full is connected, so RC_reach holds natively. The only surviving native obstruction is $H^1(Q\text{\_full})$, a holonomy/κ-type class, which obstructs RC_path, not RC_reach. The companion derivations of the Born rule settle which RC the chain consumes: reachability, not path-independence — their per-path phase is itself a holonomy, so nontrivial loop structure is the *source* of interference, and

path-independence would collapse the framework to classical. Theorem 9.11 therefore closes what the Born chain needs, and H^1 is not a threat to reconstruction but the κ/Ω closure object arriving from the reconstruction side; whether it vanishes, and whether it identifies with the Gate-3 class, are the remaining definite computations.

---

1. Introduction

The companion sequence progressively compressed the Born-rule residue. What initially appeared as several independent questions —

- Reversible Connectedness (RC),
- the Admissible Lift Property (ALP),
- the Operational Sector Principle (OSP),
- fiber–base identification,
- reversible connectedness of admissible refinements —

was subsequently shown to consist of closely linked reformulations of one underlying issue.

The danger in that is plain: a programme can manufacture the appearance of progress by reformulating a problem repeatedly without attacking it. The purpose of this paper is therefore methodological as much as mathematical. Rather than deriving yet another equivalent form of RC, we identify the formulation that possesses an *established resolution theorem* — and attack that one.

That formulation is **bracket generation**. Its merit is twofold: it carries a standard theorem (Chow–Rashevskii) on the connectedness side, and its negation carries a standard structure theorem (Frobenius, or Stefan–Sussmann in the singular case) on the disconnectedness side. So whichever way the generator computation falls, the outcome is a characterised structure rather than a restatement — which is exactly the property the previous RC-adjacent formulations lacked.

---

2. Refinement Space and Admissible Motion

Let $\mathscr{R}$ denote the admissible refinement space. Points of $\mathscr{R}$ represent admissible capacity partitions of a fixed unresolved capacity region. Let $\mathcal{A}$ denote the admissible reversible motion space, with a projection

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

associating each admissible motion configuration with the refinement it instantiates.

The admissible reversible motions define a distribution

$D \subset T\mathscr{R}$

on the tangent bundle of refinement space: $D(R) \subseteq T\_R\mathscr{R}$ is the set of directions accessible at R through admissible reversible refinement motion. The central question is whether D generates all of $T\mathscr{R}$ — and, crucially, *generates* is to be read through the Lie bracket, not pointwise span, as §4 makes precise.

A regularity note, stated now because it governs the whole analysis. Nothing so far guarantees that D has the same dimension at every point of $\mathscr{R}$. A distribution whose rank is constant is **regular**; one whose rank jumps at some points is **singular**. The admissible moves inherited from IA have no a priori reason to be regular, and the two cases are governed by different theorems (§5). We therefore keep the regularity status explicit throughout rather than assuming it silently.

---

3. Admissible Generators

We adopt the admissible motion structure inherited from IA. Each admissible infinitesimal move preserves (1) capacity, (2) admissibility, (3) reversibility, and (4) pre-commitment status.

Let

$G = \{X_1, X_2, \ldots, X_n\}$

be the complete set of admissible infinitesimal generators (vector fields on $\mathscr{R}$), and

$D = \text{span}\{X_1, \ldots, X_n\}$

the distribution they span pointwise. The key object is not D itself but the Lie algebra it generates:

$\text{Lie}(D) = \text{span}\{ X_i, [X_i, X_j], [X_i, [X_j, X_k]], \ldots \}$,

the smallest Lie algebra of vector fields containing every generator. The structure of Lie(D) — specifically, its pointwise dimension $\text{Lie}(D)(R) \subseteq T\_R\mathscr{R}$ — determines reachability.

---

4. Bracket Generation

Definition. D is *bracket-generating* at $R \in \mathscr{R}$ if $\text{Lie}(D)(R) = T\_R\mathscr{R}$. It is *globally bracket-generating* if this holds at every $R \in \mathscr{R}$.

The intuition is the control-theoretic one: the primitive admissible moves $X_1, \ldots, X_n$ may not point in every direction at a given refinement, but their iterated commutators may generate new effective directions. Bracket generation asks whether repeated admissible motion eventually reaches every admissible direction — the infinitesimal form of asking whether admissible refinements are mutually reachable.

4.1 The concrete generators: capacity transfers (established realisation)

The generic form of the admissible generators can be written down from FP1 and IA, and doing so simplifies the bracket question dramatically. Locally, an admissible refinement is a partition of the unresolved region M into parts with capacities

$R = \{C_1, \ldots, C_n\}$, $C_i = \mathrm{Vol\_op}(P_i)$, $\Sigma_i C_i = \mathrm{Vol\_op}(M)$,

so $\mathscr{R}$ is locally a **capacity simplex** with tangent space $T\_R\mathscr{R} = \{ v \in \mathbb{R}^n : \Sigma_i v_i = 0 \}$. The primitive admissible moves are **capacity transfers** between parts:

$X_{ij} = \partial/\partial C_j - \partial/\partial C_i$,

which decrease C_i and increase C_j by equal amounts, preserving ΣC_i. These are the FP1/IA-native generators: they redistribute capacity (the only datum FP1 supplies) without creating external labels (which IA excludes) and without changing total capacity (which admissibility requires).

A structural simplification follows, and it is worth stating because it collapses the general Lie-theoretic apparatus of §§4–5 to something finite. The transfer fields X_{ij} are *constant* vector fields on the simplex (their components do not depend on the C_k), so their Lie brackets vanish:

$[X_{ij}, X_{kl}] = 0$.

The commutators therefore generate no new directions: Lie(D) = D for these generators, and bracket generation reduces to *first-order* span. The question "does Lie(D) fill $T\_R\mathscr{R}$?" becomes simply "do the transfer fields span the simplex tangent space?" — which is a standard fact about graphs.

4.2 Span reduces to graph connectivity (established)

Regard the parts as vertices and the admissible transfers as edges:

vertices = parts P_i, edge i—j present iff X_{ij} is IA-admissible.

The transfer fields $\{X_{ij} : i—j \text{ admissible}\}$ span the full simplex tangent space $\{ \Sigma v_i = 0 \}$ **if and only if this graph is connected.** (This is the graph-Laplacian fact: the edge vectors $\partial_j - \partial_i$ span the hyperplane $\Sigma v_i = 0$ exactly when the kernel of the graph Laplacian is one-dimensional, i.e. the graph is connected.) Hence, for capacity-transfer generators:

connected admissible-transfer graph ⟺ D spans $T\_R\mathscr{R}$ ⟺ bracket-generating (at first order) ⟺ Branch A, disconnected admissible-transfer graph ⟺ sealed components ⟺ Branch B.

So the entire fork of §5, for these generators, is decided by one question: **is the IA-admissible transfer graph connected?** The Lie brackets do no work; connectivity is the whole content. What IA does *not* immediately settle is which pairs (i, j) are admissibly connected — that is the adjacency rule, and it is the genuine open input (§9). The next two sections establish what *can* be said about that adjacency rule from FP1 and IA alone.

---

5. The Fork Theorem

The dichotomy "Lie(D) either spans $T\mathscr{R}$ or lies in a proper sub-distribution" is trivially true and carries no content by itself. The theorem worth stating is the *consequential* one: that each side of the dichotomy yields a definite structural verdict on RC — connectedness on one side, sealed decomposition on the other. That is what the following establishes, and unlike the bare dichotomy it is not immediate: each branch invokes a non-trivial structure theorem and carries a regularity condition.

Theorem 5.1 (Fork) (conditional, as stated per branch)

Let $\mathscr{R}$ be a connected smooth manifold and D the admissible distribution of §3. Exactly one of the following holds, with the stated consequences.

Branch A — bracket-generating. If $Lie(D)(R) = T\_R\mathscr{R}$ at every $R \in \mathscr{R}$, then by the Chow–Rashevskii theorem any two points of $\mathscr{R}$ are joined by a piecewise-admissible trajectory tangent to D. Admissible reversible motion is therefore globally transitive on $\mathscr{R}$: ALP holds, and hence RC holds.

Branch B — non-bracket-generating, regular case. If Lie(D) has *constant rank* $k < \dim \mathscr{R}$, then Lie(D) is an involutive distribution of constant rank, and by the Frobenius theorem $\mathscr{R}$ foliates into leaves of dimension k. Admissible motion is confined to a single leaf; no admissible trajectory crosses between leaves. ALP fails, and hence RC fails.

Branch B′ — non-bracket-generating, singular case. If Lie(D) does *not* have constant rank (its dimension jumps at some points of $\mathscr{R}$), the Frobenius theorem does not apply and there is in general no clean foliation. By the Stefan–Sussmann orbit theorem, $\mathscr{R}$ instead partitions into *orbits* of the admissible motion — immersed submanifolds of possibly varying dimension — and admissible motion is confined to a single orbit. ALP fails on this partition, and RC fails, but the "sealed leaves" are now orbits of non-uniform dimension rather than a regular foliation.

Proof. The trichotomy A / B / B′ is exhaustive: either Lie(D) attains full rank everywhere (A), or it is proper with constant rank (B), or it is proper with non-constant rank (B′). Branch A is the Chow–Rashevskii theorem applied to the bracket-generating distribution D on the connected manifold $\mathscr{R}$. Branch B is the Frobenius theorem: a constant-rank involutive distribution is integrable, and Lie(D) is involutive by construction (it is closed under bracket). Branch B′ is the Stefan–Sussmann orbit theorem, which guarantees the orbit partition for any family of vector fields without a constant-rank hypothesis. Confinement of admissible motion to a leaf (B) or orbit (B′) is immediate: a trajectory tangent to $D \subseteq \mathrm{Lie}(D)$ cannot leave the integral manifold / orbit of Lie(D) through its starting point.

Remarks on the hypotheses

Two hypotheses are load-bearing and must not be dropped silently.

Connectedness of $\mathscr{R}$ (Branch A). Chow–Rashevskii connects points *within a connected component*. If $\mathscr{R}$ is topologically disconnected before admissible motion is considered, bracket generation connects each component internally but not across components, and "RC" must then be read as reachability within a component. Branch A as stated assumes $\mathscr{R}$ connected; if it is not, the conclusion is RC-within-components, which is strictly weaker. Whether $\mathscr{R}$ is connected as a space is itself a structural fact about the refinement geometry, logically prior to the bracket question.

Constant rank (Branch B vs B'). The clean "foliation into sealed leaves" picture of earlier RC papers is the *regular* case (Branch B). It requires Lie(D) to have constant rank — an assumption the admissible generators are not known to satisfy. If rank jumps, the correct statement is the orbit partition of Branch B′, which still seals admissible motion (RC still fails) but does not deliver uniform-dimensional leaves. The distinction matters because it is the point at which the rigour of the bracket formulation could otherwise leak: invoking Frobenius without constant rank would be an error, and we route around it explicitly via Stefan–Sussmann.

---

6. Consequences of Branch A

Suppose D is globally bracket-generating and $\mathscr{R}$ connected. By Theorem 5.1(A), admissible reversible motion is globally transitive: any admissible refinement is reachable from any other by an admissible trajectory. This is precisely the Admissible Lift Property, so

bracket generation (+ $\mathscr{R}$ connected) ⇒ ALP ⇒ RC,

the last step by the established RC ⇔ ALP analysis of the companion papers. The bath branch then follows: RC ⇒ mixing ⇒ ℓ^2 ⇒ Born-rule recovery. The quantum reconstruction chain remains viable end to end.

We note the one inherited premise explicitly: the conclusion RC (rather than RC-within-components) requires $\mathscr{R}$ connected (§5 remark). If that fails, what Branch A delivers is mixing within each component, and the cross-component structure becomes a separate question —

though one structurally similar to Branch B's sealing, now between components rather than leaves.

6.1 Two senses of RC — reachability and path-independence

A distinction must be drawn here, once and explicitly, because it governs the reading of the whole §9 analysis and resolves an otherwise-latent tension. "Reversible connectedness" can mean either of two properties, and they are not the same topologically:

- **RC_reach (reachability / transitivity).** Any admissible refinement is reachable from any other by *some* admissible trajectory. This is the property §5–6 establish from bracket generation via Chow–Rashevskii, and it is an $\mathbf{H^0}$ property: it fails exactly when $\mathscr{R}$ (or its quotient) is disconnected, so that motion is confined to separate components/leaves. Branch B's "sealed leaves" failure is the failure of RC_reach.
- **RC_path (path-independent reachability).** Not only is any refinement reachable, but the *trajectory* between two refinements is unique up to deformation — there is no holonomy around loops. This is a strictly **stronger, $\mathbf{H^1}$** property: it can fail even when RC_reach holds, on a connected space with non-contractible loops (every two points joined by a path, but the paths not all deformable into one another).

The two come apart precisely on a connected space with nontrivial H^1: there RC_reach holds (everything reachable) while RC_path fails (the route is not unique up to deformation). The torus is the standard picture — connected, fully reachable, yet $H^1 \neq 0$.

Which one the Born chain needs — resolved by the companion derivations. The chain RC ⇒ bath ⇒ mixing ⇒ ℓ^2 ⇒ Born recovery rests on RC at the first step, but on *which* RC? The companion derivations of the Born rule settle this, and in the favourable direction. Because this claim is load-bearing for the present paper, we pin it to specific results. In the kernel-uniqueness route (*The Double Square Rule*, §5.1–5.5: bilinear kernel ⇒ positive-semidefinite ⇒ factorisation-forces-rank-one (Thm 5.2) ⇒ $\varphi(P) = e^{i\theta(P)}$ (Thm 5.3) ⇒ $P = |\psi_A|^2$ (§5.5), consolidated as Thm 5.9), the Born rule is extracted by *kernel uniqueness over a path-sum* — there is no ergodic mixing step in the derivation at all, so the property consumed is that R_A is a well-defined reachable set making the path-sum $\psi_A = \Sigma_{P \in R_A} e^{i\theta(P)}$ coherent. This is **RC_reach**. The same holds in the admissibility route (*Physical Necessity of Quantum Probability Structure*, §§2–9: phase ⇒ pairwise ⇒ quadratic), which likewise builds probability from a coherent path-sum and nowhere requires route-independence. Neither route consumes RC_path.

Indeed the relationship to RC_path is the reverse of a threat, and this too is pinned. In the companion derivations the per-path phase $\theta(P)$ *is* a holonomy — defined as such in *The Double Square Rule* §3.3 ("Holonomy and the Emergence of Phase," $\theta(P) = \oint \omega$), and shown to be forced and necessarily nontrivial in *Physical Necessity* §2 and Appendix C (Layer-1 impossibility of finite holonomy; Corollary "Phase Structure," C.7). The interference cross-terms — the $2\cos(\theta(P) - \theta(P'))$ that the Born rule squares into existence — are precisely the holonomy of the loop $P \circ P'^{-1}$. Nontrivial holonomy is therefore the *source* of phase and interference, not an obstruction to them: if RC_path held (no holonomy, $\theta(P) - \theta(P') = 0$ for every same-endpoint pair), the kernel collapses to $W \equiv 1$ — exactly the $n = 0$ / decohered case that *The Double Square Rule* §5.4 and §5.7 identify as classical path-counting, *not* quantum mechanics (*Physical Necessity* §2.1, §2.4 state the same: trivial holonomy ⇒ classical, not a different quantum theory). So path-independence is not merely unneeded by the Born chain; it would *defeat* it. The chain needs reachability and wants holonomy.

This resolves the fork: the Born chain needs RC_reach, which §9 (Theorem 9.11) closes natively, so the quantum reconstruction is not hostage to H^1. We carry one honest caveat rather than collapse the fork entirely. The phase-holonomy that sources interference in the companion derivations lives on the isometry group acting on distinguishability space (*Double Square Rule* §3.1, §3.3), which is a priori a different structure from $H^1(\mathcal{Q}\_\text{full})$ (the merge/split composition-loop structure of the refinement quotient); and the companion derivations' "reachability of R_A" lives on micro-path space, a priori distinct from RC_reach on the refinement quotient. Identifying these across the two settings is the same $\mathcal{Q}\_\text{full} \simeq \Gamma\_\text{vac}$-type step flagged as open (§9.3, §9.11). What is settled, independent of that identification, is the *property* the Born chain consumes: reachability, not path-independence. The §9 results close the H^0 (reachability) obstruction natively; the surviving H^1 residue obstructs only RC_path, which the Born chain does not require — and which, were it trivial, would return a classical rather than quantum framework.

(Note: the "bath ⇒ mixing ⇒ ℓ^2" phrasing of the chain derives from Born Rule as Entropic Unfolding *(Part I); the present resolution does not depend on that route, since the kernel-uniqueness route above reaches the Born rule with no mixing step and consumes only RC_reach. The pinned citations are to the two routes that do not pass through mixing, which is what makes the "reachability, not path-independence" conclusion robust to how the mixing route is formulated.)*

Theorem 6.2 (Born reconstruction requires RC_reach but not RC_path). *Assume the companion derivations' premises: (i) a reachable outcome-set R_A; (ii) coherent path amplitudes $\psi\_A = \Sigma\_{P\in R\_A} e^{i\theta(P)}$ with $\theta(P)$ the per-path holonomy; (iii) the kernel-uniqueness derivation of the Born rule (rank-one kernel ⇒ $P(A) = |\psi\_A|^2$). Then RC_reach is sufficient for the reconstruction, and RC_path is not required. Moreover, if RC_path holds — i.e. every same-endpoint pair of paths is route-equivalent, so all loop holonomies $\theta(P) - \theta(P')$ vanish — the interference kernel collapses to $W \equiv 1$ and the reconstruction returns the classical path-counting limit, not quantum mechanics.*

Proof. (Sufficiency of RC_reach.) Reachability makes R_A a well-defined set and ψ_A a coherent sum; the kernel-uniqueness derivation (Double Square Rule §5.1–5.5, Thm 5.9) then yields $P(A) = |\psi\_A|^2$ with no further hypothesis on the routes. No step invokes route-uniqueness, so RC_path is not consumed. (Collapse under RC_path.) If all loop holonomies vanish, then for any $P, P' \in R\_A$, $\theta(P) - \theta(P') = 0$, so the kernel $W(P, P') = e^{i(\theta(P)-\theta(P'))} \equiv 1$; this is the $n = 0$ case of Double Square Rule §5.4, giving $P(A) \propto |R\_A|$ — classical path-counting (§5.7, decohered limit). Hence RC_path is not merely unneeded; imposing it removes interference.

The collapse clause is scoped to the *phase-holonomy* that enters the kernel — the holonomy of the isometry group on distinguishability space (Double Square Rule §3.1, §3.3) — which is the object the reconstruction consumes. Its identification with $H^1(\mathcal{Q}\_full)$ on the refinement quotient remains the open $\mathcal{Q}\_full \simeq \Gamma\_vac$ step (§9.3); Theorem 6.2 does not assume that identification, and is a statement about the reconstruction's own holonomy.

Then why care about RC_path at all? A reader may reasonably ask: if Born reconstruction needs only RC_reach, and Theorem 6.2 shows RC_path is not part of it, why does the paper continue to track RC_path? Because RC_path has *changed jobs*. It is no longer a condition on reconstruction — Theorem 6.2 removes it from that role — but it is exactly the condition that decides the *closure* question. RC_path holds iff $H^1(\mathcal{Q}\_full) = 0$ (Theorem 12.1), and $H^1(\mathcal{Q}\_full)$ is the κ-type Gate-3 closure class (§9.8, §9.11). So RC_path migrated from the reconstruction side of the programme to the closure-topology side: its failure does not threaten the Born rule (which is secured by RC_reach), but it *is* the Gate-3 residue. Tracking RC_path is therefore tracking whether the sevenfold closure sector leaves a topological trace — a separate and independently interesting question from "does quantum probability follow," now cleanly disentangled from it. The two senses of RC partition the programme: RC_reach is the reconstruction question (closed, Thm 9.11), RC_path is the closure question (open, = the κ computation).

---

7. Consequences of Branch B / B′

Suppose D is not bracket-generating. Then admissible motion is confined — to a leaf (regular case, B) or an orbit (singular case, B′) — and no admissible trajectory crosses between them. ALP fails; RC fails. This is exactly the sealed-refinement branch identified in earlier papers.

The payoff for the broader programme is that the operational sector boundaries are *identified* here, not posited: the sectors are the leaves (or orbits) of the non-bracket-generating distribution. No separate sector-boundary programme is required — the sector structure is the geometric manifestation of non-bracket-generation, and the question "can a sector boundary exist without an operational witness?" becomes "does the leaf/orbit carry an operational invariant?", answerable with the same machinery. The regular/singular distinction also tells the sector analysis what kind of object it is dealing with: uniform-dimensional sealed sectors (B) versus sectors of varying dimension with possible singular strata (B′).

This subsumes the proposed "Operationally Invisible Sector Boundaries II" line: it is not a separate paper but the Branch-B reading of this one. The sealed leaves *are* the sector boundaries, and characterising them is the Branch-B corollary, not an independent programme.

---

8. The No Internal Seal Lemma

The fork theorem says *if* the transfer graph is disconnected, sealed leaves follow. The natural next question is whether disconnection can arise from FP1 and IA at all — that is, whether a seal separating admissible refinements can be built from the primitives, or whether it must be imported as extra structure. This section constrains the answer by exhausting what a seal could be made of.

8.1 The seal trichotomy

Suppose, for contradiction, that $\mathscr{R}$ has sealed components: a function $S : \mathscr{R} \to \Lambda$ with $S(R) = S(R')$ iff R, R′ lie in the same component, constant along every admissible trajectory. For the seal to be real, S must be determined by the primitive structure. FP1 supplies only the region M, its operational volume Vol_op(M), the resolution scale Δ_op, the distinguishability dimension d_op, and the capacity bound $|\Sigma(M)| \leq \text{Vol\_op}(M)/\Delta\_\text{op}^{\{d\_op\}}$. It does **not** supply a pre-labelled outcome set, a conserved sector charge, a fixed partition menu, or a rule forbidding capacity exchange between refinements. Given this, S can be of only three kinds:

1. an **external label** on refinements;
2. a **capacity-derived invariant**, computed from the C_i;
3. a **dynamical conserved invariant**, accumulated along admissible motion.

The trichotomy is exhaustive: S must be supplied from outside (1), computed from the only datum FP1 gives (2), or conserved by the dynamics (3). There is no fourth source, because FP1's ontology contains nothing else.

8.2 Case 1 — external label: excluded by IA (proven)

An external sector label is exactly what IA forbids: refinements are internally individuated by the region's own capacity structure, not selected from an externally supplied label set. The No Pre-Individuation result is explicit that FP1 attaches capacity to a region, not to pre-existing labelled outcomes. So S cannot be an external label.

8.3 Case 2 — capacity-derived invariant: constant, hence no seal (proven)

Suppose S is built from the capacities alone. To seal components, S must be invariant under every admissible transfer X_{ij}. But the transfers can move capacity between any IA-adjacent parts while preserving only the total ΣC_i; a quantity invariant under all of them can depend on the capacities only through ΣC_i. And ΣC_i = Vol_op(M) is **the same for every refinement of M**. Therefore any transfer-invariant capacity function is globally constant on $\mathscr{R}$ and separates nothing.

This case is closed by genuine impossibility, not parsimony: a capacity-derived seal is not merely unmotivated, it is mathematically constant. This is the strongest step of the lemma — it eliminates the most plausible candidate seal outright.

8.4 Case 3 — dynamical conserved invariant: cannot come from FP1, but not closed (relocated, not eliminated)

The only remaining possibility is that S is a quantity conserved along every admissible reversible path — $S(R(\tau)) = \text{const}$. Here we must be exact about what is and is not established.

First, what such an S *is*: a conserved quantity whose level sets are the sealed components is precisely the sealed-leaf structure of Branch B. So "S is a dynamical seal," "the transfer graph is disconnected by a conservation law," and "RC fails" are the *same statement*. Case 3 is therefore not a third independent threat — it is RC-failure itself, named. The lemma's real content is that RC-failure has exactly *one* possible form, the dynamical seal, the other two having been excluded.

Second, what FP1 says about it: FP1 is not a dynamical axiom. It states capacity and finite distinguishability; it supplies no conservation law over refinement sectors. So a dynamical seal **cannot be derived from FP1** — if one exists, it is carried by the admissible-motion dynamics (the adjacency rule of §4.2), which is specified beyond FP1, not by FP1 itself.

Third — and this is the boundary the lemma must not overstep — "FP1 does not supply it" is **not** "it does not exist." The dynamics of admissible motion is real structure, and whether *that* structure conserves a sector charge is the residual open question. Case 3 is therefore *relocated*, not closed: pushed out of FP1 and into the adjacency rule, which is exactly the step-3 input the bracket reduction already identified as gating (§9). The lemma narrows where a seal could live — only in the dynamics, nowhere in FP1 — without proving the dynamics is seal-free.

8.5 The parsimony step for the residual gap

One move remains, and it is defeasible, so we mark it as such. Even Case 3 is not *logically impossible*: a theory FP1 + IA + (dynamical seal) is consistent. But the seal it posits is operationally idle relative to FP1 data — by construction it is a conserved charge over sectors, and Case 2 showed capacity cannot witness it, so FP1's data cannot register it at all. The closing move is therefore the parsimony one: *an operationally idle conserved charge is not posited unless forced, and nothing in FP1 + IA forces it.* Under this principle RC holds. We are explicit that this is a parsimony judgement, not a proof of inconsistency — the same defeasible-but-principled footing on which analogous residual gaps elsewhere in the programme are handled. It raises the burden onto anyone positing a seal to exhibit what operational work it does; it does not prove no consistent theory contains one.

The "operationally idle" claim can be sharpened from an impression into a proposition, which strengthens the parsimony step by making precise exactly what the seal cannot do.

Operational Detectability Proposition (proven, relative to FP1). *Let R, R' be two refinements differing only by the value of a Case-3 seal Q (same component of the admissible-transfer structure, distinguished only by the conserved charge). Then R and R' have identical capacities (Q is not capacity-derived, §8.3), identical admissible observations (any FP1 observable is a* *function of the capacity/distinguishability data, which agree), and identical transfer dynamics (the admissible moves act the same way on both, since they depend on capacities, not on Q). Therefore no FP1 observable distinguishes R from R': the seal has no witness among the observables the primitive theory defines.*

This moves the claim from "the seal seems idle" to "the seal has no observable witness in the primitive theory" — a definite statement about FP1's observables, not an impression. We mark its boundary honestly, because it is the point a hostile reader will press: the proposition proves no-witness *relative to FP1's observables*, not no-witness *simpliciter*. Many real physical structures are operationally hidden relative to some restricted set of observables and detectable relative to a richer set; the proposition does not claim the seal is undetectable in principle, only that *FP1's own data cannot reach it*. That is exactly what the parsimony step needs: positing the seal adds structure the theory's own observables cannot register, so the theory gains a degree of freedom it can never use. It does not need, and does not assert, that no enriched theory could detect such a charge — only that within FP1 + IA it is witnessless, hence unmotivated.

8.6 What the lemma establishes

No Internal Seal Lemma. Given FP1 and IA, an admissible refinement space admits no nontrivial sealed components except those induced by a dynamical conserved invariant carried by the admissible-motion dynamics — and such an invariant is not contained in FP1, must be added as extra structure, and is operationally idle relative to FP1 data.

Equivalently, and at full honesty: **FP1 + IA do not by themselves force RC, but they force any RC failure to be caused by an extra dynamical conserved seal not contained in FP1.** The negative half is proven — there is no FP1/IA-native obstruction to admissible paths between refinements (Cases 1–2 closed outright). The positive half — that admissible reversible paths actually exist — requires the dynamics to carry no sector charge, which is the residual question, addressable by parsimony (§8.5) but not closed deductively.

This is the sharpest available statement of the RC residue: not "why RC?" but "does the admissible-motion dynamics conserve a sector charge, or not?" — a property of the adjacency rule, computable once that rule is written down.

---

9. The Dynamical Seal No-Go (Partial) — What a Conserved Seal Must Be

Case 3 of the No Internal Seal Lemma left the dynamical seal *possible but uncharacterised*. We now constrain its form, proving what such a seal must look like, and then — carefully separating the proven part from the conjectural part — trace where that constraint leads.

9.1 A conserved seal must be deformation-invariant (proven)

Let $Q : \mathscr{R} \to \Lambda$ be a dynamical conserved seal: $Q(R(\tau))$ constant along every admissible reversible trajectory. By the No Internal Seal Lemma, Q is neither an external label (Case 1) nor capacity-derived (Case 2). We show it must be invariant under all metric deformation.

The admissible trajectories are generated by the capacity transfers X_{ij} (§4.1), and a conserved Q is constant along each. But the transfers move *every* metric quantity continuously: within a refinement's combinatorial type, any capacity assignment can be deformed into any other by a sequence of transfers (this is the connectivity-of-the-simplex content of §4.2, now used internally rather than across types). So Q must take the same value on all capacity assignments related by transfer — it cannot depend on the metric data (the actual values of the C_i) at all. By the same argument that made Case 2's capacity-seal constant, Q is invariant under continuous metric deformation of the refinement.

Proposition 9.1. *Any admissible conserved seal Q is metric-deformation-invariant: it is constant on each set of refinements related by admissible capacity transfer, and therefore cannot be a metric quantity. Q can depend only on the non-metric structure of the refinement — its combinatorial / relational skeleton.*

This is a real narrowing and it is proven: the seal, if it exists, lives on the discrete skeleton of refinement space, not on its metric fibres. Everything the transfers can move is disqualified; only deformation-invariant structure survives.

9.1A Killing the non-topological discrete candidates — conditional on merge/split

Proposition 9.1 leaves a discrete candidate set: part-count, ordering, lattice position, and genuinely topological invariants. We now show that all but the last die — *under one explicit premise about the admissible-motion structure.*

The premise concerns whether admissible motion can change the *combinatorial type* of a refinement, not merely redistribute capacity within a fixed type. A fixed-arity capacity transfer X_{ij} moves capacity between existing parts and leaves the number of parts n unchanged. To connect refinements of different part-counts, one needs **merge/split moves**: an admissible move that combines two parts into one ($n \to n-1$) or divides one into two ($n-1 \to n$).

Premise (No Privileged Carving). Admissible refinement motion includes merge/split moves. The motivation is internal to IA: IA individuates refinements internally from one unresolved region and privileges no particular carving. A carving into n parts and a carving into n−1 parts are both admissible refinements of the *same* region; forbidding passage between them would privilege the part-count as primitive structure — precisely the externally-fixed structure IA denies. So IA's no-privileged-carving spirit motivates merge/split admissibility. Whether IA's *formal* statement delivers it is a separate question, addressed below.

Proposition 9.1A (conditional). *Under the No Privileged Carving premise, any admissible conserved seal Q is invariant under change of part-count, ordering, and description. Consequently part-count seals, ordering seals, and lattice-position seals are all excluded: a merge move changes the part-count and the lattice position, and admissible relabelling changes the ordering, so none of these can be conserved along admissible motion. The only surviving candidates are invariants of the refinement's intrinsic relational structure — independent of how it is carved, counted, or ordered.*

Proof. Q is conserved along every admissible trajectory. Under the premise, merge/split trajectories are admissible, so Q is constant across them: a refinement and its one-merge image have equal Q. Part-count and lattice position both change under a merge, so a Q distinguishing them could not be constant — contradiction; they are not seals. Ordering and description-dependent quantities change under admissible relabelling (itself a within-type admissible symmetry), so by the same argument they are not seals. What remains invariant under merges, splits, and relabellings is exactly the intrinsic relational structure of the refinement, independent of its presentation.

What this does to Conjecture 9.2. Under the premise, the candidate set collapses from "any deformation-invariant discrete structure" to "intrinsic relational invariants only" — and the non-topological discrete candidates that made Conjecture 9.2 a genuine open question (part-count, lattice position) are gone. The conjecture then narrows: the obvious non-topological discrete candidates (part-count, lattice position) are gone, so a seal that is neither metric (9.1) nor presentational (9.1A) is *pushed toward* the topological. We are careful not to overstate this — it is a narrowing, not an elimination. "Intrinsic relational invariant being exhausted by cohomology" is the same open identification Conjecture 9.2 names (and §9.2 states the gap explicitly); 9.1A closes the obvious non-topological escapes but does not by itself prove the survivor is cohomological. The route is much narrower; it is not yet closed.

The load-bearing premise — and why it is not independent. Proposition 9.1A is stated as conditional on No Privileged Carving (that IA admits merge/split). It would seem this leaves an independent open premise: does IA's formal statement permit arity-changing moves or only fixed-arity carvings? The next subsection shows the premise is *not* independent — denying it (insisting on fixed arity) requires positing that part-count is conserved, which is itself a seal, and which runs the §8 trichotomy and lands in the same Case-3 dynamical-charge branch as every other seal. So "merge/split admissible" is not a separate assumption but the statement "no arity-conservation seal is added," subject to the parsimony verdict already established. We make this precise now.

9.1B No Fixed-Arity Seal — the premise is an instance of the trichotomy

There appear to be two candidate quotient complexes, differing only in whether part-count is quotiented away:

Q_fixed = $\mathscr{R}$ / (relabelling + capacity deformation + fixed-arity presentation changes), *preserves part-count;* Q_full = $\mathscr{R}$ / (relabelling + capacity deformation + presentation changes + merge/split), *quotients part-count away.*

The whole merge/split question is whether part-count is an admissible invariant before commitment — i.e. whether Q_fixed is a legitimate *native* quotient rivalling Q_full. We show it is not: fixed arity is not native to FP1 + IA, and Q_fixed is merely Q_full equipped with an added dynamical arity-seal — a special case of Branch B, not a rival.

Theorem 9.1B (No Fixed-Arity Seal). *Given FP1 and IA, fixed arity is not a primitive admissible invariant. If part-count is conserved, it is a Case-3 dynamical seal — non-native to FP1, not capacity-derived, and subject to the §8.5 parsimony verdict.*

Proof. Suppose fixed arity is conserved: admissible motion cannot cross between $\mathscr{R}_n$ (refinements with n parts) and $\mathscr{R}_{n+1}$. Then n functions as a conserved seal Q(R) = n. Run it through the trichotomy of §8:

- *External label?* No — n being a primitive sector label is exactly the externally-supplied individuation IA excludes. FP1 supplies M, Vol_op(M), Δ_op, d_op and a capacity bound; it supplies no "n-sector" label. A fixed-arity quotient quietly reintroduces the pre-commitment label "this refinement belongs to the n-sector," which is what IA was designed to exclude.
- *Capacity-derived?* No — by the §8.3 argument: total capacity ΣC_i = Vol_op(M) is unchanged under subdivision, so a 3-part and a 4-part carving of M have identical capacity data; n is not a function of the capacities.
- *Presentation-invariant under coarsening?* No, trivially — coarsening is *defined* as the move that changes n; n is precisely the quantity refinement/coarsening alters.

So n is not external, not capacity-derived, and not coarsening-invariant. It can survive only as a Case-3 dynamical conserved invariant — an added arity-conservation law not contained in FP1.

Consequence — Q_fixed is not a rival native quotient. Fixed arity is therefore not a second native option competing with Q_full. It *is* Q_full plus a dynamical arity-seal — an instance of Branch B, the sealed branch, with the seal being specifically a conserved part-count. The earlier conditional "under No Privileged Carving" can now be replaced by the stronger and more honest "unless an arity-conservation seal is added": FP1 + IA ⇒ Q_full, modulo the same residual possibility that governs every seal — an extra Case-3 dynamical charge, here a conserved n. By §8.5, such a charge is operationally idle relative to FP1 data and parsimony declines to posit it; we are explicit, as always, that this is a parsimony verdict, not a proof that an arity-conserving theory is inconsistent.

Net effect. The part-count seal does not survive as an *independent* non-topological seal. It is itself just an extra dynamical seal — the same object §8 already classified. So the canonical quotient for the paper is Q_full, and the fixed-arity version is relegated to "what happens if the dynamics adds an arity-conservation seal." The merge/split premise of §9.1A is thereby absorbed into the seal trichotomy rather than standing outside it as an independent assumption: Proposition

9.1A holds *unless an arity-conservation seal is added*, which is the §8.5 parsimony condition applied to one named candidate.

9.2 What the non-metric skeleton contains — and the gap

It is tempting to conclude immediately that Q is therefore *topological* — a homology, cohomology, or winding class — and to read the seal as a Gate-3-type invariant. We must not take that step as proven, because there is a genuine gap, and naming it is what keeps this section honest.

"Metric-deformation-invariant" is **not** synonymous with "topological invariant." Proposition 9.1 pins Q to the structure that survives capacity transfer *within a combinatorial refinement type*, and to whatever distinguishes the types themselves. The type-distinguishing data need not, a priori, be topological in the homological sense: it could be merely combinatorial (the number of parts n), order-theoretic (the position of the carving in the refinement lattice), or some other discrete invariant of the carving. Propositions 9.1A–9.1B close these specific escapes: part-count, ordering, and lattice position die under merge/split, and §9.1B shows that *denying* merge/split is not a free alternative but requires positing a conserved part-count — itself a Case-3 dynamical seal subject to the parsimony verdict. So the canonical reading eliminates these candidates, with their survival relegated to the arity-sealed Branch-B case. What remains is the residual identification that the surviving *intrinsic relational* invariants are exactly the cohomological ones; "deformation-invariant and presentation-invariant" establishes Q is an intrinsic discrete invariant, and the final step to "cohomological" is what Conjecture 9.2 still names. The step from "intrinsic relational invariant of refinement structure" to "topological invariant of the Gate-3 type" is the remaining identification, narrowed by 9.1A–9.1B but not eliminated.

We therefore state the topological reading as a conjecture with its gap explicit, not as a theorem:

Conjecture 9.2 (Topological Seal). The deformation-invariant charges on refinement space are the cohomology classes of its combinatorial skeleton; equivalently, every admissible conserved seal Q is represented by a nontrivial topological invariant of the refinement complex. Under this conjecture, RC failure ⇒ nontrivial topology.

The gap between Proposition 9.1 (proven) and Conjecture 9.2 (open) is precisely whether the deformation-invariants of refinement space are *exhausted by* its cohomology, or whether non-cohomological discrete invariants (part-counts, lattice positions) can also seal. That is a structural fact about $\mathscr{R}$, not settled here.

9.3 The conditional bridge to Gate-3

The reason Conjecture 9.2 matters is that, *if* it holds, the RC residue and the Gate-3 closure sector become the same kind of object — and the paper's earlier deliberate move away from Gate-3 reverses into a bridge, built without assuming it. The logic is:

RC failure ⇒ conserved seal (No Internal Seal Lemma, §8) ⇒ metric-deformation-invariant charge (Proposition 9.1, proven) ⇒ [Conjecture 9.2] nontrivial topological invariant of the refinement complex ⇒ [if that complex is Γ_vac] a class in the cohomology of Γ_vac — a candidate Gate-3 residue.

Two named identifications stand between the proven part and the Gate-3 conclusion: Conjecture 9.2 (deformation-invariants are cohomological, not merely combinatorial), and the identification of the refinement complex with Γ_vac (that RC's refinement skeleton and Gate-3's vacuum complex are the same object). Neither is established here. With both, RC-failure and the Gate-3 residue are literally one object; with only Proposition 9.1, they share the weaker property that both obstructions are deformation-invariant discrete charges.

9.4 The same gap, twice — a structural observation

It is worth recording that the gap in Conjecture 9.2 is the *same shape* as the central open question of the companion Continuum-Lift paper. There, the question was whether the Gate-3 $\mathbb{Z}_7$ is genuine homological torsion in $H_1(\Gamma\_vac; \mathbb{Z})$ or merely a coefficient choice — i.e. whether a discrete structure is genuinely topological or only notationally so. Here, the question is whether refinement space's deformation-invariant charges are genuinely cohomological or merely combinatorial — i.e. whether a discrete structure is genuinely topological or only apparently so. Both papers reach a wall of the same form: *metric/continuous data cannot carry the discrete seal (proven in each), and whether the surviving discrete data is truly topological is the open question (conjectured in each).*

Indeed the proven halves are structurally identical: this paper's Case 2 (capacity-derived seals are constant, §8.3) is the analogue of the Continuum-Lift paper's Proposition 4.1 (continuous-holonomy admits no canonical discrete class). Two independently-developed lines hitting the same wall — *only topology can seal, and whether the candidate is genuinely topological is open* — is evidence that the deformation-invariant-charge story points at something real. It is not proof; "the same gap appears twice" supports taking the bridge seriously without building it. But it does mean the two programmes' residual questions are the same kind of question, and resolving the topological-versus-combinatorial status of refinement space's invariants would bear on both.

9.5 Classification of admissible conserved seals

The preceding subsections eliminate candidate seals one at a time. We now assemble the eliminations into a classification — their cumulative force is a single structural statement, not a list of rejections. The assembled result is the **Minimal Residue Theorem**, stated and proved at §9.9; this subsection sets up the classification it rests on and the chain it anchors. The chain organising all of §9:

1. **Eliminate** all local candidates — metric (§9.1), label (§8.2), ordering (§9.5 below), part-count/arity (Theorem 9.1B), local combinatorial (§9.8), history (§9.9). Each is killed by a specific admissible move.

2. **Descend**: every surviving obstruction factors through the quotient *Q* — the Minimal Residue Theorem (Theorem 9.9, universal property).
3. **H^0**: under full merge/split, *Q*_full is connected, so the component-seal vanishes and RC_reach holds (Theorem 9.11).
4. **H^1 remains**: the only surviving native seal is an H^1 holonomy — the κ-type, obstructing RC_path (§9.8, §9.11).

Steps 1–3 are proven; step 4 is the sharply-localised open residue. The rest of this subsection proves step 1's invariance requirements; §9.9 states and proves the descent (step 2); §9.8 and §9.11 supply steps 3–4.

Theorem 9.5A (Invariance requirements). *Any admissible conserved seal Q must be:*

1. *Metric-invariant* — invariant under all admissible capacity transfers, hence not a metric quantity (Proposition 9.1, proven unconditionally).
2. *IA-invariant* — independent of outcome labels, external indexing, and any externally-supplied structure, since IA excludes these (No Internal Seal Lemma, Case 1, proven).
3. *Presentation-invariant* — unchanged under admissible relabelling and, under the No Privileged Carving premise, admissible regrouping (merge/split); otherwise Q measures bookkeeping rather than admissible structure (Proposition 9.1A, conditional on merge/split for the regrouping clause).

Corollary (eliminated candidates). The following are excluded as seals:

- *Capacity moments* (mean, variance, any function of the C_i) — metric; excluded by clause 1, unconditionally.
- *Outcome / part count* — changes under admissible regrouping; excluded by clause 3 *under No Privileged Carving*. (Survives as a candidate if IA is fixed-arity only — see below.)
- *Ordering* — changes under admissible relabelling; excluded by clause 3, unconditionally (relabelling is a within-type symmetry IA grants regardless of arity).
- *Lattice position* — presentation-dependent, changes under regrouping; excluded by clause 3 *under No Privileged Carving*.

Proposition 9.5B (Descent to the quotient complex). *After eliminating metric, label, and presentation data, any surviving seal Q descends to an invariant of the quotient*

Q = (refinements) / (admissible equivalence),

and is therefore an invariant of the resulting combinatorial complex. By §9.1B the canonical quotient is *Q*_full (merge/split included): fixed arity is not a rival native quotient but *Q*_full plus a dynamical arity-seal. So the canonical reading is that admissible equivalence includes regrouping, outcome-count is quotiented out, and Q is an invariant of the fully-reduced *Q*_full — *unless* an arity-conservation seal is added, in which case the relevant complex is the finer *Q*_fixed and outcome-count survives as exactly that added seal. Either way Q descends to a quotient; what an added arity-seal changes is only *which* quotient, and per §9.1B that addition is itself a Case-3 dynamical charge subject to the §8.5 parsimony verdict, not an independent native option.

Note what the classification did and did not use. It eliminated by appeal only to metric-invariance, IA, and presentation-invariance. It never invoked Gate-3, never invoked κ, never invoked cohomology. The conclusion — *a surviving seal is an invariant of the quotient refinement complex* — was reached by elimination alone.

9.6 The final fork

Theorem 9.6 (RC ⇔ quotient invariant). *Relative to the quotient complex Q of Proposition 9.5B, exactly one holds:*

Branch A. *Q possesses no nontrivial admissible invariant. Then every admissible conserved seal is constant, the transfer graph is connected, and RC holds.*

Branch B. *Q possesses a nontrivial admissible invariant. Then that invariant is a candidate seal, the transfer graph may be disconnected by it, and RC fails — with the seal necessarily a deformation- and presentation-invariant invariant of Q.*

Equivalently, modulo the specification of Q:

RC ⇔ the quotient refinement complex Q has no nontrivial admissible invariant.

This is a classification result, not a localisation. The earlier sections established that *if* a seal exists it is not metric, not a label, not presentational; §9.5 collected these into the positive statement that it must be an invariant of Q; and §9.6 makes RC equivalent to the absence of such an invariant. The residue is no longer "is there a mysterious charge?" but "does this specific combinatorial complex carry a nontrivial invariant?" — a question about a definite object.

We flag the one premise-dependence honestly: the biconditional is relative to Q, and the canonical Q is Q_full (§9.1B). Adding an arity-conservation seal would shift the relevant complex to the finer Q_fixed — but by §9.1B that addition is a Case-3 dynamical charge, not a co-equal native quotient, so the biconditional's canonical form is against Q_full, with Q_fixed appearing only in the arity-sealed branch of Branch B.

9.7 Why the Gate-3 bridge is now a corollary, not a premise

The classification makes the Gate-3 connection unavoidable in the right way — as the answer to a question the classification forces, rather than as an imported assumption. Having established (§9.6) that a surviving seal must be a nontrivial invariant of the quotient complex Q, the natural question is simply: *what invariants does a combinatorial complex possess?* The standard answer is its homotopy and (co)homology invariants — H_0, H_1, cohomology classes, holonomy. So the seal, if it exists, is one of these — and $\kappa \in H^1(\Gamma\_vac; \mathbb{Z}_7)$ enters precisely as the Gate-3 instance of "a nontrivial cohomology class of the refinement complex," *if* $\mathcal{Q}$ is Γ_vac (the identification of §9.3).

This is the disciplined form of the bridge. We did not assume RC and Gate-3 are related and then look for a connection. We eliminated until only quotient-complex invariants survived, then observed that the invariants of a complex are cohomological, then noted κ is the Gate-3 cohomology class. Each step is a corollary of the prior one; the only assumption is the identification $\mathcal{Q}$ = Γ_vac, which is named and not smuggled. Conjecture 9.2 is now exactly the statement that the *natural* (cohomological) invariants of $\mathcal{Q}$ are *all* its admissible invariants — the last gap between the classification and the topological reading, isolated to a single question about $\mathcal{Q}$.

9.8 Seal topology: components and loops

The classification (§9.5–9.6) says a seal must be an invariant of $\mathcal{Q}$ but leaves "invariant of a complex" informal. We now make it precise, eliminate the last non-topological candidate (local combinatorial structure), and identify exactly which topological invariants a seal can be. The result completes the elimination programme: metric, label, presentation, and now local-combinatorial data are all excluded, leaving precisely the global topological invariants H^0 and H^1.

The locally-constant core. A conserved seal Q is invariant under every admissible move, so it cannot change under any infinitesimal admissible motion: Q is *locally constant* on $\mathcal{Q}$. A locally constant function on a connected space is globally constant. Therefore:

Proposition 9.8A (H^0 core, proven). *A nontrivial admissible conserved seal forces $\mathcal{Q}$ to be disconnected: it is a nonconstant locally-constant function, hence determined by the components of $\mathcal{Q}$, i.e. by $H^0(\mathcal{Q})$. If $\mathcal{Q}$ is connected, every seal is constant and RC holds.*

This is the rigorous core, and note where it lands: on **H^0 (components)**, not H^1 (loops). The primary, simplest seal detects *disconnection* — admissible motion failing to bridge between components — which is a π_0/H^0 phenomenon and needs no loop at all. The sealed-leaves picture of Branch B (§7) is exactly this: the leaves are the components of $\mathcal{Q}$, and "sealed" means $H^0(\mathcal{Q})$ is nontrivial.

Two seal types, distinguished. There is a second, finer seal mechanism, and it must be separated from the first because it is the one that connects to Gate-3:

- **H^0 component-seal** — sealed components; admissible motion cannot cross between them. Detected by disconnection of $\mathcal{Q}$. The primary seal; loopless.
- **H^1 holonomy-seal** — within a *connected* component, a quantity conserved around non-contractible loops (a holonomy). Requires the component to be connected but to have nontrivial $\pi_1(\mathcal{Q})$; the seal is then an H^1/holonomy invariant. This is the κ type.

These are genuinely different: a component-seal needs disconnection (H^0); a holonomy-seal needs connectedness-with-loops (H^1). An earlier framing that demanded "every seal detects a non-contractible loop" would be too strong — it would miss the H^0 component-seal, which is loopless and is in fact the simpler and more primary form of RC-failure. The correct statement covers both.

Eliminating local combinatorial structure. The last non-topological candidate is local combinatorial data — local incidence, local part-structure, anything defined on a neighbourhood within a component. This dies by the same argument that killed ordering and lattice position (§9.5): admissible motion within a component can vary local combinatorial data continuously, so a locally-constant Q cannot depend on it. What admissible motion *cannot* vary is the global structure — how many components there are, how loops close within them, and (potentially) higher-order global features. So local combinatorial structure is excluded, and the survivors are global invariants of Q's topology.

Proposition 9.8B (Cornering). *Every admissible conserved seal is forced onto the global topology of Q: metric, label, presentation, and local-combinatorial candidates are all eliminated, so a surviving seal must be a global topological invariant. The two mechanisms identified here are H^0 (sealed components) and, within a component, H^1/holonomy (sealed loops). These are the only presently-identified seal mechanisms; whether they exhaust the global invariants of Q is the open question of the next paragraph.*

This converts Conjecture 9.2 from "is the seal topological at all?" into the much sharper "the seal is global-topological, and is it captured by the identified H^0/H^1 mechanisms or does Q carry a higher one?" The elimination of *non-topological* candidates is, at this level, complete: there is no surviving metric, label, presentational, or local-combinatorial place for the seal to hide. What is not yet complete is the classification of the *topological* survivors by degree.

The one residual gap, now sharp. What remains genuinely open is whether H^0 and H^1 *exhaust* the admissible invariants of Q. Higher cohomology — H^2 and above — could in principle carry a seal, if Q's structure admits nontrivial higher classes that are also admissible-invariant. The locally-constant argument forces *component*-seals into H^0 and the holonomy argument forces *loop*-seals into H^1, but neither rules out a higher-degree invariant acting as a seal through some more elaborate mechanism — so H^0 and H^1 are the identified mechanisms, not a proven-exhaustive list. The honest residue is therefore: the seal is forced onto the global topology of Q, with H^0 and H^1 the only presently-identified mechanisms, and whether a higher-degree mechanism exists is a question about a specific complex's H^{2+} — no longer about whether the seal is topological at all. That is a far narrower and more answerable gap than "is the seal topological?": it presupposes topology and asks only about degree.

The sharpened Gate-3 bridge. With the two seal types separated, the Gate-3 connection becomes precise. $\kappa \in H^1(\Gamma\_vac; \mathbb{Z}_7)$ is an H^1 class, so it is specifically the *holonomy-seal* type (§9.7) — the within-a-connected-component, loop-based seal — not the H^0 component-seal. So *if* $Q = \Gamma\_vac$ and the RC-failure is of the holonomy type (not the component type), the seal is exactly a class in $H^1(\Gamma\_vac; \cdot)$, and κ is its Gate-3 instance. This is sharper than "seal ⇒ topology ⇒ κ": it specifies *which* topological seal type κ corresponds to (H^1 holonomy, connected component with loops) and distinguishes it from the H^0 component-seal that would seal RC without any loop and would *not* be a κ-type object. The bridge to Gate-3 is therefore to the holonomy branch of the seal classification specifically.

9.9 The Minimal Residue Theorem

This is the headline result the elimination chain of §9.5 culminates in — the descent that makes "global invariant of the quotient" rigorous (step 2 of the §9.5 chain). It is stated and proved here, where the proof lives; §9.5 introduces it and the surrounding chain.

Setup. Let $\sim_a$ be admissible refinement equivalence on $\mathscr{R}$, generated by the admissible moves: capacity-preserving transfer, admissible relabelling, admissible presentation change, and — *if admissible* — refinement/coarsening (merge/split). Define the quotient refinement complex $\mathcal{Q} := \mathscr{R} / \sim_a$ with quotient map $\pi_a : \mathscr{R} \to \mathcal{Q}$.

Theorem 9.9 (Minimal Residue). *After quotienting away metric structure, labels, ordering, part-count, local combinatorics, and history, every surviving admissible conserved seal factors through the quotient: there is a unique $q : \mathcal{Q} \to \Lambda$ with $Q = q \circ \pi_a$, and if Q is nontrivial then q is a nontrivial invariant of $\mathcal{Q}$. Equivalently — every nontrivial admissible seal is a nontrivial global invariant of the quotient refinement complex; nothing local, metric, or presentational survives.*

Proof. A seal is invariant under every admissible transformation, so if $R \sim_a R'$ then $Q(R) = Q(R')$ — Q is constant on equivalence classes. By the universal property of the quotient, Q factors uniquely as $q \circ \pi_a$. Nontriviality of Q transfers to q. This step is unconditional: it is the universal property, and holds for whatever $\sim_a$ admissibility actually is. The content the theorem inherits from §§8–9.8 is *which* data has been quotiented into $\sim_a$ — that the eliminated candidates are genuinely eliminated — so that the residue q is a *global* invariant and not a repackaged local one.

Which candidates are eliminated, and by which move. The descent disposes of every candidate that two $\sim_a$-equivalent refinements can differ by — and it is worth recording the precise correspondence, because each non-topological class is killed by a *specific* admissible move:

- *capacity transfer* kills metric data (capacity values, moments) — unconditional (§8.3, §9.1).
- *relabelling* kills ordering and external indexing — unconditional (IA, §9.5).
- *coarsening (merge/split)* kills local combinatorial structure — part-count, local valence, adjacency, branching pattern, tree depth — **conditional on coarsening being admissible** (No Privileged Carving, §9.1A).
- *refinement history independence* is then a corollary, not a separate axiom: two histories reaching the same refinement are presentation-equivalent, so a presentation-invariant Q cannot distinguish them.

The first two are unconditional; the third is the load-bearing one. Local combinatorial features are precisely those a coarsening move can change — a single merge alters part-count, valence, adjacency, and tree depth at once — so *if* coarsening is admissible, no local combinatorial quantity is conserved, and local-combinatorial seals die. This is the Local Combinatorial Seal

Elimination, and it is exactly as conditional as merge/split admissibility: it is the coarsening clause of $\sim_a$ doing the work.

What the premise fixes. The descent (Theorem 9.9) is unconditional, but *which* complex Q the seal descends to depends on whether $\sim_a$ includes coarsening. By §9.1B this is settled canonically in favour of Q_full:

- *Canonical case — coarsening admissible* (Q_full): $\sim_a$ is the full equivalence, Q is maximally quotiented, and metric / label / presentation / local-combinatorial / history are all eliminated. The seal is a global invariant of the fully-reduced complex — a "global residue" in the strict sense. This is canonical because, per §9.1B, denying coarsening requires positing a conserved part-count, which is itself a Case-3 dynamical seal rather than a native feature.
- *Arity-sealed case* (Q_fixed): if a dynamical arity-conservation seal is added, $\sim_a$ is finer, Q is larger, and local combinatorial invariants (part-count and its relatives) survive — but as the content of that added seal, not as independent native invariants. This is a special case of Branch B (the sealed branch), not a rival native quotient.

The honest stopping point. Theorem 9.9 proves *seal ⇒ nontrivial invariant of Q*. It does **not** prove *seal ⇒ cohomology* or *seal ⇒ κ* — that would be too fast, and the descent does not deliver it. What kinds of invariant Q can carry is the separate question answered in §9.8: a seal is locally constant (hence H^0, components) or a holonomy (hence H^1, loops), with higher cohomology the residual open degree-question. So §9.9 is the *descent* — it puts the seal onto Q — and §9.8 is the *identification* — it says which invariants of Q a seal is known to be. Together: every nontrivial seal is a global invariant of Q (§9.9), forced onto its topology, with H^0 and H^1 the identified mechanisms and a higher-degree mechanism the one open possibility (§9.8). Topology enters not because Gate-3 was assumed, but because everything metric, labelled, presentational, local-combinatorial, and historical has been quotiented away, and what survives a maximal quotient is by definition global.

9.10 The complete elimination and the residual target

It is worth assembling the eliminations into a single statement, because their cumulative force is the paper's main result and is easy to lose across nine subsections. A surviving admissible conserved seal has been excluded from every non-global hiding place:

Candidate seal	Eliminated by	Status
External label / outcome index	IA (No Internal Seal Lemma, Case 1)	excluded, unconditional
Capacity metric (total, moments)	capacity-transfer invariance (Case 2, Prop 9.1)	excluded, unconditional
Ordering	admissible relabelling (§9.5)	excluded, unconditional
Local combinatorial structure	coarsening (§9.9), via the arity argument	excluded — see part-count row

Candidate seal	Eliminated by	Status
Part count / fixed arity	Theorem 9.1B: requires a Case-3 dynamical seal	excluded as native; survives only as added seal
Refinement history	presentation-invariance (corollary, §9.9)	excluded, unconditional

The part-count row is what closes the local-combinatorial row: local combinatorial structure is eliminated *by coarsening*, and Theorem 9.1B shows that denying coarsening is not a free native alternative but requires positing a conserved part-count — itself a Case-3 dynamical seal. So the canonical quotient is Q_full, and every non-global candidate is gone.

Strongest honest conclusion. What remains is exactly a *global invariant of Q_full*, and by §9.8 that invariant is H^0 (sealed components) or H^1 (holonomy) unless Q_full carries seal-bearing higher cohomology. So:

FP1 + IA do not force RC, but they force any RC failure to be an added conserved global invariant of the full quotient refinement complex Q_full — an H^0 or H^1 class (or, in the one open case, a higher-cohomology class) — not contained in FP1 and operationally idle relative to FP1 data. Under parsimony (§8.5): no added invariant ⇒ RC.

We keep "global invariant of Q_full, which is H^0/H^1 modulo higher cohomology" rather than the looser "global topological/dynamical invariant," because the topological identification is *earned* for H^0 and H^1 (§9.8) and *open* for higher degree; collapsing them would re-open the gap §9.8 isolated.

The residual target, as a burden. The result reframes RC-failure as a claim someone must now *exhibit*, against a very narrow specification. To exhibit a surviving seal — and thereby fail RC — one must produce a quantity on Q_full that is simultaneously:

1. not metric (survives all capacity transfers),
2. not label-based (admissible under IA),
3. not fixed-arity (not a conserved part-count — Theorem 9.1B),
4. not local-combinatorial (survives coarsening),
5. conserved by the admissible dynamics, and
6. operationally meaningful relative to FP1 data.

Requirements 1–4 force the quantity onto the global topology of Q_full; 5 makes it a genuine conservation law of the dynamics; and 6 is the one §8.5 flags as the sticking point — by construction such a seal is idle relative to FP1's data, so 6 is hard to meet, which is precisely why parsimony declines to posit it. This is a very narrow target, and naming it is the point: the paper does not prove no such quantity exists, but it has reduced "why RC?" to "exhibit a quantity meeting all six, on a specific complex" — which is a constructive challenge against a sharp specification, not an open-ended doubt.

This is as much as can be closed without explicitly constructing $\mathcal{Q}$_full and computing its invariants. That construction — and the higher-cohomology exhaustion question of §9.8 — is the remaining work, and it is now a definite computation on a definite complex rather than a conceptual question about whether RC "should" hold.

9.11 The minimal full quotient — connectedness, and why not contractibility

The full quotient can be partly computed directly, and the computation closes the H^0 (component) seal natively while sharpening the residual question to H^1.

Connectedness (proven). Under full admissible merge/split equivalence, any finite partition reduces to the trivial one-part partition {M} by repeated merging, and any partition is reached from {M} by repeated splitting. So for any two refinements R, R′ there is an admissible path R → {M} → R′, hence $R \sim_a R'$ at the level of connectivity.

Theorem 9.11 (Minimal Full Quotient — H^0). *Under FP1, IA, and full admissible merge/split equivalence, $\mathcal{Q}$_full is connected:*

$H^0(\mathcal{Q}\_\text{full}) = \mathbb{Z}$ (one component).

No component-seal exists: the H^0 form of RC-failure is eliminated natively. Equivalently, the part-count sectors $\mathscr{R}_n$ are all merged into one — which is Theorem 9.1B seen geometrically, since fixed arity was the only thing that could have disconnected $\mathcal{Q}$_full, and full merge/split removes it.

Why this does not give contractibility — and why that matters. It is tempting to conclude $\mathcal{Q}$_full ≃ *point*, hence $H^k(\mathcal{Q}\_\text{full}) = 0$ for all $k \geq 1$, hence no seal at all and RC closes outright. This does **not** follow, and the gap is exactly the one that carries the Gate-3 connection. The merge-to-{M} argument proves *path-connectedness* (H^0), not contractibility. A connected space can still have nontrivial loops: the circle is connected, every two points joined by a path, yet $H^1(S^1) = \mathbb{Z}$. The argument shows any two refinements are *connected* by a merge/split path; it does not show the *different* merge/split paths between them are deformable into one another. If two essentially distinct merge/split routes from R to R′ cannot be continuously deformed into each other, $\mathcal{Q}$_full carries a nontrivial loop and $H^1(\mathcal{Q}\_\text{full}) \neq 0$ — even though it is connected.

Concretely, the merge/split moves compose like the morphisms of a refinement category, and $\mathcal{Q}$_full is the geometric realisation of that composition structure. Connectivity follows because the moves link every partition to {M}; but H^1 is governed by the *relations among the moves* — whether merge-then-split and split-then-merge cycles commute, or whether there are composition cycles that do not bound. Whether such cycles exist is a fact about how the admissible moves compose — the same adjacency/composition structure that is the paper's standing open input — not something the connectivity argument settles.

The residual obstruction is exactly $H^1(Q\_full)$ — and it is the κ-type seal, not an RC_reach failure. Combining with §9.8 and the disambiguation of §6.1: under full merge/split, Q_full is connected (Theorem 9.11), so the H^0 component-seal is dead and **RC_reach holds natively** — every refinement is reachable from every other, exactly as Proposition 9.8A states ("if Q is connected, RC holds"). What a nontrivial H^1 obstructs is *not* reachability but path-independence: on a connected Q_full, $H^1 \neq 0$ means non-contractible loops, so routes between refinements exist (RC_reach) but are not unique up to deformation (RC_path fails). H^1 is the holonomy/κ seal type (§9.7–9.8), and it is the obstruction to RC_path, leaving RC_reach intact. Therefore:

Under the native FP1/IA full quotient, RC_reach holds (Theorem 9.11: Q_full connected). The only surviving native obstruction is $H^1(Q\_full)$, which obstructs RC_path, not RC_reach — and that H^1, if present, *is* a κ-type closure class.

This resolves what would otherwise be a contradiction with Proposition 9.8A. That proposition says connected ⇒ RC (meaning RC_reach), and Theorem 9.11 gives connected, so RC_reach holds — full stop. The H^1 residue does not reopen RC_reach; it is a distinct, strictly finer obstruction (RC_path), the path-independence of reconstruction. Calling it "RC-failure" without qualification would conflate the two senses §6.1 separates and would falsely contradict 9.8A. Stated correctly: the H^0 form of RC (reachability) is *closed* natively; the H^1 form (path-independence) is the open residue, and it is the κ-object.

This is a sharper statement than "Q_full = point, RC closes," and it is the honest one. It does not overclaim $H^1 = 0$ (which the connectivity argument cannot deliver and which would, if asserted, simultaneously close RC_path *and* sever the Gate-3 bridge the rest of §9 builds). Instead it identifies the residual RC_path obstruction with a single cohomology group, $H^1(Q\_full)$, and identifies that group with the Gate-3 κ-object. The two programmes' open questions, which §9.4 noted were the *same shape*, are now under full merge/split *literally the same group*: H^1 of the refinement quotient. That is the unification made precise.

What this means for the Born chain — see §6.1. The companion derivations settle which RC the Born chain consumes: reachability, not path-independence (pinned to *Double Square Rule* §3.3, §5.1–5.5 and *Physical Necessity* §2 in §6.1). So Theorem 9.11 closes what the quantum reconstruction needs, the H^1 residue is not a threat to Born recovery but the κ/Ω closure object, and — since the per-path phase is itself a holonomy — the residue is structurally the *same kind of object* the Born derivation is built on, which strengthens rather than threatens the Gate-3 bridge. The one open caveat (identifying the companion phase-holonomy with $H^1(Q\_full)$) is the $Q\_full \simeq \Gamma\_vac$ step of §9.3.

Contractibility would require the extra fact. If the merge/split composition structure has *no* nontrivial cycles — if every merge/split diagram commutes up to deformation — then $H^1(Q\_full) = 0$, Q_full is contractible, RC_path holds as well, and there is no κ-class. That is a genuine possibility, but it requires establishing the no-composition-cycles fact about the moves; it does not follow from connectivity. (Note that since the Born chain needs only RC_reach, contractibility is not required for reconstruction — it would simply mean the κ-class is trivial.) So contractibility is not claimed — it is flagged as the additional structural input that would close

H^1 (and hence RC_path), exactly parallel to the higher-cohomology question of §9.8 (which closes H^{2+}) and the adjacency question of §11.

The caveat, attached to the right group. The earlier caveat — restricted merge/split gives nontrivial Q_full and an added seal — now attaches precisely to H^0 / RC_reach: restricting merge/split *disconnects* Q_full, reintroducing a component-seal (an RC_reach failure), which is exactly the arity-seal of Theorem 9.1B. So the seal structure splits cleanly by degree and by RC-sense: a *restricted* merge/split adds an H^0 component-seal (RC_reach failure, the arity-seal, non-native by 9.1B); and *even with full* merge/split, a nontrivial composition-loop structure would carry an H^1 holonomy-seal (RC_path failure, the κ-type, the one genuinely open native possibility).

9.11A Necessary conditions for an H^1 obstruction

Theorem 9.11 leaves H^1(Q_full) open. A referee is entitled to ask the symmetric question to "why should H^1 vanish?" — namely *why should H^1 be nontrivial at all?* We do not answer it (that would require the composition law of the moves, which we do not have), but we can state what a nontrivial H^1 *requires*, and observe that FP1 and IA push against each requirement. The effect is to make a nontrivial H^1 look exceptional rather than equally likely — without proving it absent.

A nontrivial class in H^1(Q_full) is a non-contractible loop in the merge/split composition structure: a cycle of admissible moves that returns to its starting refinement but cannot be deformed to the trivial loop. Such a cycle requires at least one of the following:

1. **Noncommuting merge/split operations** — two moves whose order matters up to deformation, so that a merge-then-split cycle does not bound. If all admissible merges and splits commute up to deformation, every composition cycle is fillable and $H^1 = 0$.
2. **Path-dependent refinement histories** — a refinement reachable by two routes that are not deformation-equivalent, so the difference of routes is a nontrivial loop. If refinement is history-independent (the same refinement is "the same" however reached), routes are interchangeable and the loops bound.
3. **An obstruction to contracting admissible cycles** — some topological feature of the move-structure that prevents a closed sequence of moves from being shrunk, even when the individual moves commute locally.

Now the pressure FP1 and IA exert on each:

- *Against (1):* the concrete moves are capacity transfers and merge/splits on a capacity simplex (§4.1). Capacity transfers are constant commuting vector fields ($[X_{ij}, X_{kl}] = 0$, §4.1); merge and split are mutually inverse operations on partitions. Mutually inverse, capacity-additive operations have no obvious source of non-commutation — the natural composition is abelian. Non-commutation would require extra structure on the moves beyond capacity-additivity, which FP1 does not supply.
- *Against (2):* IA's internal individuation and §9.5's presentation-invariance already quotient out refinement history — a refinement is individuated by its capacity structure,

not by the route that produced it. So in Q_full, history-dependence has been deliberately quotiented away; route-difference loops are, to the extent the quotient is faithful, already collapsed. Residual path-dependence would have to survive the very quotient designed to remove it.
- *Against (3):* a contraction obstruction is a genuine topological feature, and this is the one FP1/IA do *not* obviously suppress — it is exactly the open possibility. But note it cannot come from (1) or (2) if those are closed; it would have to be an intrinsic topological feature of the fully-quotiented complex, which is precisely the $H^1 = \kappa$ question, not a separate mechanism.

What this establishes, and what it does not. Conditions (1) and (2) are the mechanisms FP1/IA push against: commuting capacity-additive moves and a quotient that removes history both work to fill composition cycles. Condition (3) is the residue — an intrinsic topological contraction-obstruction of Q_full — and it is exactly $H^1(Q\_full)$ = the κ-class, not a new mechanism. So the analysis does not prove $H^1 = 0$; it shows that the *non-topological* routes to a nontrivial H^1 (non-commuting moves, path-dependence) are the ones FP1/IA resist, leaving only the *intrinsically topological* route — which is the κ question itself. A nontrivial H^1 therefore cannot arise from bookkeeping artifacts of the moves; it must be genuine topology of the quotient, if it exists at all. This makes a nontrivial H^1 exceptional — it has to be real topology, not an artifact — without proving it absent. We state plainly: this is a narrowing of *how* H^1 could be nontrivial, not a proof that it is trivial; proving $H^1 = 0$ would require the composition law we do not have, and we do not claim it.

9.11B Local commutation does not decide global H^1

The arguments of §9.11A must not be read as proving $H^1(Q\_full) = 0$. Local commutation of admissible moves is *not* sufficient to eliminate global first cohomology: the torus has locally commuting directions everywhere and still carries nontrivial H^1. So "the moves commute up to local deformation" does not close the question — a globally nontrivial loop can survive even when every local diagram commutes.

The point of §9.11A is narrower and, made precise, stronger. It is not that H^1 vanishes; it is that a nontrivial H^1, *if present*, cannot be blamed on representational artifacts — ordering, arity, history, or local non-commutation — because those have been quotiented away or shown FP1/IA-resisted. Therefore any surviving H^1-class must be a genuine global feature of the fully-reduced quotient itself. The residue is no longer "perhaps the moves were described badly"; it is exactly

$H^1(Q\_full) \neq 0$

as an intrinsic topological fact about the quotient, or not at all. This is the strongest honest narrowing available: the paper does not prove the class vanishes, but it proves that if it exists, it is not a representational artifact — it is real topology of Q_full. The distinction matters because it forecloses the cheap dismissal ("you just described the moves badly") and the cheap closure ("the moves commute, so $H^1 = 0$") in the same stroke: neither the artifact-objection nor the local-commutation-shortcut decides H^1, and what remains is a genuine computation of the quotient's first cohomology.

9.11C Why the residue is in degree 1 — the lowest-degree theorem

The paper repeatedly locates κ in H^1. That placement should be earned, not inherited from the fact that Gate-3 currently points there. The following theorem earns it: it shows degree 1 is the *first* degree that can host a surviving seal, because degree 0 is closed.

Theorem 9.11C (Lowest-Degree Residue). *Assume (i) $H^0(Q\_full)$ is closed — Q_full connected, no component-seal (Theorem 9.11); (ii) admissible seals are forced onto the global topology of Q_full (Prop 9.8B); (iii) any such seal is, by the seal-type classification, a cohomological invariant of Q_full. Then the first (lowest-degree) admissible residual obstruction is H^1. Degree 0 provably carries none (it is closed), so the lowest degree that can host a surviving seal is 1, and loop holonomy realises it. Higher cohomology (H^{2+}) may also carry seals, but cannot be the first surviving obstruction, since $1 < 2$.*

Proof. By (i), $H^0(Q\_full) = \mathbb{Z}$ and the only locally-constant invariant is the constant — no degree-0 seal survives (Prop 9.8A). By (ii)–(iii), a surviving seal is a cohomological invariant of degree ≥ 1. The smallest such degree is 1, and it is non-vacuous: a loop holonomy on a connected complex is a genuine admissible H^1 invariant (§9.8, the holonomy-seal type), so degree 1 is realised as a candidate rather than excluded. Ordering the admissible degrees, degree 0 is closed and degree 1 is the first that can carry a seal; any H^{2+} class, if present, has strictly higher degree and so cannot be the lowest-degree (first) obstruction.

What this does and does not settle. It settles *where the leading residue lives*: degree 1, not 0, and ahead of any higher degree — so $\kappa \in H^1$ is not an assumption imported from Gate-3 but the lowest available slot once H^0 is closed. It does **not** settle that H^{2+} is empty; that is the §9.8 exhaustion question, left open. The two are consistent: H^1 is the *first* surviving obstruction (proven, given H^0 closed), and whether additional seals sit in higher degree is a separate, higher-degree computation (open). So the placement of κ in degree 1 is now licensed by the closure of degree 0, independent of the Gate-3 pointer — which is what makes the §9.7 bridge a corollary rather than a premise.

---

10. Reduction of the RC Problem

The analysis changes the *kind* of question RC is, in two stages. The bracket fork (§5) showed RC is not a global topological survey of $\mathscr{R}$ but a local generator question. The concrete realisation (§4.1–4.2) then showed the generators are capacity transfers, for which the brackets vanish and the question collapses to graph connectivity. And the No Internal Seal Lemma (§8) showed that the only way the graph can be disconnected is a dynamical conserved seal not contained in FP1. Putting these together, the residual problem is:

1. The admissible generators are the capacity transfers $X_{ij} = \partial/\partial C_j - \partial/\partial C_i$ (known generically, §4.1).
2. Their span fills T_R$\mathscr{R}$ iff the IA-admissible transfer graph is connected (§4.2); the Lie brackets add nothing.
3. The graph can fail to be connected only via a dynamical conserved sector charge carried by the adjacency rule — not from FP1, which excludes label-seals (Case 1) and capacity-seals (Case 2) outright (§8).

So the entire RC branch reduces to a single question about the **IA-admissible adjacency rule**: does it connect the parts, or does the admissible-motion dynamics conserve a sector charge that disconnects them? Connected adjacency ⇒ Branch A ⇒ RC ⇒ ℓ^2/Born. A conserved disconnecting charge ⇒ Branch B ⇒ sealed leaves ⇒ ledger branch — and by §8 such a charge is extra structure beyond FP1, operationally idle relative to FP1 data. This is a sharper reduction than the bracket fork alone delivered: not "compute Lie(D)" (now trivial — the brackets vanish), but "is the adjacency rule connected, equivalently does the dynamics carry a non-FP1 sector charge?"

10A. The Next Computation: Constructing $\mathcal{Q}$_full

The preceding sections reduce RC to a definite cohomology computation, but they do not yet construct the complex on which that computation is to be performed. With the reduction in hand, the construction programme is now explicit. One must specify:

1. **Objects** — admissible refinements of a fixed unresolved region M.
2. **Morphisms** — admissible capacity transfers, relabellings, presentation changes, and merge/split moves.
3. **Relations** — equivalences identifying paths that differ only by admissible presentation, by the order of independent moves, or by refinement history.
4. **Geometric realisation** — the quotient complex $\mathcal{Q}$_full generated by those objects, morphisms, and relations (the nerve of the refinement category modulo the relations).
5. **Cohomology** — compute $H^0(\mathcal{Q}\text{\_full})$, $H^1(\mathcal{Q}\text{\_full})$, and check whether higher classes can carry admissible seals.

This paper proves the first cohomology computation already (Theorem 9.11):

$H^0(\mathcal{Q}\text{\_full}) = \mathbb{Z}$ (connected; no native component-seal, so RC_reach holds — §6.1, §9.11).

The remaining native obstruction is therefore the next computation:

$H^1(Q\_full)$,

the holonomy/κ-type class, which obstructs RC_path (path-independence), not RC_reach (reachability). This makes the next step unavoidable rather than speculative: the reduction has converted "does RC hold?" into "construct Q_full from the specification above and compute its first cohomology," and the H^0 rung of that computation is already done. The construction is a definite object built from data the programme either has (the moves) or has localised as its one open input (the composition law / adjacency rule of the moves, §9.11, §11). What this paper does not do — and flags as the work that follows — is carry out steps 2–5 for the *actual* admissible-motion structure; that requires the composition law written down, at which point $H^1(Q\_full)$ is a computation rather than a question.

---

11. What Remains Unknown

The paper establishes the *form* of the generators (capacity transfers), proves two of the three seal-cases impossible, and proves any surviving seal must be metric-deformation-invariant — but it does not establish the adjacency rule, and therefore does not establish RC. The open problem is now a single, sharply-stated question, with a characterisation of what its negative answer would have to look like.

Open Problem. Determine the IA-admissible adjacency rule on refinement parts: which pairs (i, j) admit an admissible capacity transfer X_{ij}. Equivalently — by §8 — determine whether the admissible-motion dynamics conserves any sector charge beyond FP1. A connected adjacency rule (no such charge) gives Branch A and RC; a disconnecting conserved charge gives Branch B and the sealed-leaf branch. By the No Internal Seal Lemma this charge, if it exists, is not contained in FP1 and is operationally idle relative to FP1 data (parsimony, §8.5, favours its absence defeasibly); and by Proposition 9.1 it must be metric-deformation-invariant — a discrete invariant of the refinement skeleton, not a metric quantity.

Subsidiary open questions, now sharply posed by §9:

- *Arity-conservation seal (Theorems 9.1B, 9.11).* The canonical quotient is Q_full, and under full merge/split it is *connected* (Theorem 9.11): the H^0 component-seal is closed natively, since fixed arity — the only thing that could disconnect Q_full — is not a rival native quotient but a Case-3 dynamical seal (Theorem 9.1B). So the residual question is not "does IA admit merge/split?" but "does the dynamics add an arity-conservation seal?", which restricting merge/split would require; parsimony (§8.5) declines it, leaving Q_full connected.
- *The H^1 residue — the κ question (§9.8, §9.11).* With H^0 closed under full merge/split, RC_reach holds natively (Q_full connected); the *only* surviving native seal is an H^1 holonomy — a non-contractible loop in the merge/split composition structure — which is exactly the κ-type seal and obstructs RC_path, not RC_reach (§6.1, §9.11). The companion derivations consume reachability, not path-independence (their phase is a

holonomy, so loop structure is the *source* of interference), so this residue does not threaten the Born chain; it is the κ/Ω closure object. The open question is whether the merge/split moves have nontrivial composition cycles ($H^1 \neq 0$, κ present) or commute up to deformation ($H^1 = 0$, $\mathcal{Q}$_full contractible). This is a definite computation on the composition structure, not a conceptual doubt.
- *Higher cohomology (§9.8).* Whether H^0 and H^1 exhaust the seal-bearing invariants, or whether H^{2+} can also seal, remains open — a question about a specific complex's higher cohomology, presupposing topology and asking only about degree.
- *The Γ_vac identification (§9.3).* Is the combinatorial skeleton of refinement space the same complex as the Gate-3 vacuum complex Γ_vac? Under full merge/split this becomes: is $H^1(\mathcal{Q}$_full) the same group as H^1(Γ_vac)? If so, the RC residue and the Gate-3 residue are literally one cohomology class — the unification made precise (§9.11).

Two prior structural facts remain presupposed and worth confirming independently: that $\mathscr{R}$ is connected as a space (required for Branch A to give RC rather than RC-within-components, §5), and the regularity status of the distribution (which selects regular sealed leaves B from singular orbits B′, should Branch B obtain).

---

12. Conclusion

The whole argument compresses into a single hierarchy, which we state before the closing summary.

Theorem 12.1 (Hierarchy). *Under FP1 and IA, with $\mathcal{Q}$_full the full admissible quotient:*

RC_reach ⇔ $H^0(\mathcal{Q}$_full$) = \mathbb{Z}$ ($\mathcal{Q}$_full connected), and RC_path ⇔ $H^1(\mathcal{Q}$_full$; G) = 0$ (given connectedness; G the holonomy coefficient group, $\mathbb{Z}_7$ for the κ object).

Consequently: the quantum reconstruction requires the first condition only (Theorem 6.2), and Theorem 9.11 establishes it natively; the Gate-3 closure question concerns the second, the H^1 residue being the κ-type class. The chain FP1+IA ⇒ No Internal Seal (§8) ⇒ Minimal Residue (Thm 9.9) ⇒ $H^0(\mathcal{Q}$_full$) = \mathbb{Z}$ (Thm 9.11, RC_reach closed) ⇒ only H^1 survives ⇒ Gate-3 candidate (§9.8, §9.11) is therefore a descent through cohomological degree: degree 0 closed, degree 1 the residue, degree ≥ 2 the §9.8 open question.

Proof. The first biconditional is Theorem 9.11 with its converse: a locally-constant seal is nonconstant iff $\mathcal{Q}$_full is disconnected (Prop 9.8A), and RC_reach is reachability between all refinements, which holds iff $\mathcal{Q}$_full is connected, i.e. $H^0 = \mathbb{Z}$. The second: on a connected complex, admissible motion is route-independent (RC_path) iff the holonomy representation of every loop is trivial; for holonomy valued in the coefficient group G this is exactly $H^1(\mathcal{Q}$_full; G) = 0 — the holonomy obstruction (§9.8, §9.11), with $G = \mathbb{Z}_7$ for the Gate-3 κ class. (The "trivial-vs-nontrivial" content is coefficient-independent; naming G makes the boxed biconditional exact rather than schematic.) The "given connectedness" qualifier is supplied by the first biconditional under full merge/split (Thm 9.11). The reconstruction clause is Theorem 6.2.

This single statement carries the paper's architecture: the two senses of RC are the two lowest cohomology groups of one complex, reconstruction needs only the degree-0 one, and Gate-3 lives in degree 1.

The previous RC literature produced several equivalent formulations of one residual question. This paper attacks the single formulation possessing established resolution theorems on *both* sides of its fork — bracket generation, with Chow–Rashevskii on the connected side and Frobenius / Stefan–Sussmann on the sealed side — and then sharpens it four times. First, the admissible generators are concretely the capacity transfers, for which the Lie brackets vanish and bracket generation collapses to connectivity of the admissible-transfer graph. Second, the No Internal Seal Lemma proves that graph can be disconnected only by a dynamical conserved seal absent from FP1: label-seals are excluded by IA, capacity-seals are mathematically constant. Third, any such seal is proven metric-deformation-invariant, presentation-invariant, and (via the seal classification) forced onto the global topology of the quotient complex $\mathcal{Q}$_full, with the non-topological candidates — metric, label, ordering, part-count, local-combinatorial — all eliminated. Fourth, the full quotient is computed far enough to show $\mathcal{Q}$_full is *connected* (Theorem 9.11), closing the H^0 component-seal natively: the only thing that could disconnect it, fixed arity, is not a rival native quotient but a Case-3 dynamical seal.

The residue is therefore as small as the present analysis can make it, and it has acquired a definite shape and a definite degree — once the two senses of RC are separated (§6.1). With H^0 closed under full merge/split, $\mathcal{Q}$_full is connected and **RC_reach holds natively** (Theorem 9.11, consistent with Proposition 9.8A: connected ⇒ reachability). The *only* surviving native seal is an H^1 holonomy — a non-contractible loop in the merge/split composition structure — and it obstructs **RC_path** (path-independence), not RC_reach. So the terminus is sharp: **reachability is closed natively; the open residue is $H^1(\mathcal{Q}$_full), the obstruction to path-independence, which is the κ-class.** If the merge/split moves commute up to deformation, H^1 vanishes, $\mathcal{Q}$_full is contractible, and RC_path holds too; if they carry a nontrivial composition cycle, that cycle is an H^1 holonomy-seal and RC_path fails. This is no longer "why RC?" or "is there a hidden charge?" but a definite computation of one cohomology group of one complex.

The residue does not threaten the Born chain: the companion derivations consume reachability, not path-independence (Theorem 6.2, pinned in §6.1, proved minimally in Appendix A), so Theorem 9.11 closes what reconstruction needs and the H^1 residue is the κ/Ω closure object, not a Born-chain obstruction — structurally the same loop-holonomy on which the Born rule is built. The cross-setting identification (companion phase-holonomy with $H^1(\mathcal{Q}$_full)) is the open $\mathcal{Q}$_full ≃ Γ_vac step; the property the Born chain consumes is settled, and it is the one Theorem 9.11 delivers.

The negative half of RC is proven down through H^0; the H^1 question is open but definite, and the positive default rests on parsimony. That is the honest terminus on RC alone. But the H^1 identification also delivers the more ambitious finding, now stated at the level it is established: the surviving H^1 holonomy-seal *is* the κ-type closure class (§9.8), so if the refinement quotient

$\mathcal{Q}$_full and the Gate-3 vacuum complex Γ_vac share their first cohomology, then the RC residue and the Gate-3 residue are *literally the same group* — H^1 of the refinement quotient. The paper does not prove that identification; it proves the rungs beneath it (the seal is forced onto global topology; under full merge/split it is specifically H^1; H^1 is the κ-type) and names the one identification ($\mathcal{Q}$_full ≃ Γ_vac in first cohomology) that would complete it. What §9.4 noted as the *same shape* of open question across the two programmes is, under full merge/split, sharpened to the *same cohomology group* — so the quantum-reconstruction residue and the closure-topology residue are not merely analogous but candidates to be one and the same object, pending a single computation.

References (companion papers)

The Born-chain resolution of §6.1 is pinned to the following companion derivations in the programme; the section numbers cited there refer to these documents.

- **The Double Square Rule: A Derivation of Quantum Probability from Discrete Informational Geometry** (Part II of the two-part Born-rule programme). Kernel-uniqueness derivation: §3.1, §3.3 (holonomy and the emergence of phase, $\theta(P) = \oint \omega$); §5.1–5.5 (bilinear kernel → positive-semidefinite → factorisation-forces-rank-one, Thm 5.2 → $\varphi(P) = e^{i\theta(P)}$, Thm 5.3 → $P = |\psi_A|^2$, §5.5); §5.4 and §5.7 (the trivial-holonomy / decohered limit is classical path-counting); Thm 5.9 (consolidated statement).
- **Physical Necessity of Quantum Probability Structure** (the assumption-justification companion). Admissibility route: §2 and Appendix C (Layer-1 impossibility of finite holonomy; phase forced and necessarily nontrivial; Corollary "Phase Structure," C.7); §2.1, §2.4 (trivial holonomy ⇒ classical, not a distinct quantum theory).
- **Born Rule as Entropic Unfolding** (Part I). Source of the "bath ⇒ mixing ⇒ ℓ^2" phrasing of the Born chain; the §6.1 resolution does **not** depend on this route, since the kernel-uniqueness route above reaches the Born rule with no mixing step.

Appendix A: Why H^1 Does Not Threaten Quantum Reconstruction

This appendix makes the paper self-contained on the one external dependency that the terminus rests on (§6.1, Theorem 6.2): that the Born reconstruction consumes reachability, not path-independence. The minimal derivation is short enough to give in full.

The reconstruction. For a macro-outcome A with reachable reversible-path set R_A, the companion derivations build the amplitude as a coherent sum of per-path phases,

$\psi_A = \Sigma_{P \in R_A} e^{i\theta(P)}$,

and the kernel-uniqueness argument (rank-one positive-semidefinite kernel forced by factorisation) yields

$P(A) = |\psi_A|^2 = \Sigma_{P,P' \in R_A} e^{i(\theta(P) - \theta(P'))}$.

The key identity. The summand depends only on the difference $\theta(P) - \theta(P')$. This difference is a *loop quantity*: it is the holonomy of the closed loop $P \circ P'^{-1}$ (go out along P, return along P′ reversed). The phase $\theta(P)$ itself is defined as a holonomy (a connection integral / sum of edge phases along P); the observable content is the relative holonomy between paths.

The two limits. Everything about whether the result is quantum or classical turns on this loop quantity:

- *Loops present (holonomy nontrivial):* $\theta(P) - \theta(P') \neq 0$ for some pairs, so the kernel carries genuine phase, the cross-terms $2\cos(\theta(P) - \theta(P'))$ survive, and P(A) exhibits interference. This is quantum mechanics.
- *No loops (holonomy trivial — RC_path):* $\theta(P) - \theta(P') = 0$ for all same-endpoint pairs, so the kernel collapses to $W \equiv 1$, every cross-term is +1, and $P(A) \propto |R_A|$ — i.e. uniform classical path-counting after normalisation, $P(A) = |R_A| / \Sigma_i |R_{A_i}|$, the decohered/diagonal limit.

The conclusion. Reachability of R_A is what the path-sum needs to be defined (RC_reach). The interference that distinguishes quantum from classical is *produced by* nontrivial loop holonomy — so path-independence (RC_path, no holonomy) is not a prerequisite of the reconstruction but its negation in the quantum regime. A nontrivial H^1-type loop structure is therefore the source of the very phenomenon the Born rule encodes, not an obstruction to deriving it. This is why the surviving $H^1(\mathcal{Q}_{full})$ residue, far from threatening reconstruction, is the same kind of object — loop holonomy → phase — on which the reconstruction is built; whether the reconstruction's phase-holonomy and $H^1(\mathcal{Q}_{full})$ are literally the same group is the open $\mathcal{Q}_{full} \simeq \Gamma_{vac}$ identification (§9.3), not a question that bears on whether H^1 threatens the Born chain. It does not.