Detecting the Gate-3 Closure Class

Primitive-Fact Holonomy, Merge–Split Motion, and the Reversible-Connectedness Residue

Keith Taylor — VERSF Theoretical Physics Programme

General Reader Summary

The claim this paper brings within reach

Every genuine new fact reality commits would leave its mark on a single, unerasable tally deep in the substrate — not a separate memory of each fact, but one quantised total that no admissible motion can smooth away.

The sentence carries two honesties. First, the paper does not prove the tally is ever nonzero — it proves the question now hangs on one yes-or-no, whether a single reversible refinement loop is, after transport, the very loop reality draws around a committed fact, with nowhere else for the answer to hide. If it is, the tally is real, permanent, and the very object the closure side of the programme found independently; if it is not, every loop smooths to nothing and reality keeps no such record. Second, the tally is a *total*, not an archive: it counts in sevens, so aligned commitments can cancel, and beneath the net the framework keeps no log of which facts came in which order. What would be permanent is one small running count of genuine commitment — not a chronicle of events, and not lasting *through* time so much as standing outside it, since it is this unerasability that lets there be a fixed past at all.

Before reality commits to a definite fact, there are many admissible ways to refine the same unresolved situation. The companion paper, *The Admissible Lift Theorem*, established the simplest fact about this collection of refinements: once admissible merging and splitting are allowed, they all belong to a single connected structure. No admissible refinement is walled off in an isolated sector; any admissible refinement can reach any other.

That closes the crudest possible obstruction — but connectedness is not the end of topology. A space can be connected and still contain loops. A traveller can move between any two points on a circle, yet the circle still carries a loop that cannot be shrunk to a point. The question of this paper is whether anything analogous survives in admissible refinement space: not whether you can *get* from one refinement to another, but whether the *different routes* between them are equivalent. And — as will matter shortly — whether any loop that does survive is the very loop the closure side of the programme has already found.

The paper's main result is not a computation. It is a statement about where the answer can possibly live. Earlier work eliminated every local candidate for a hidden obstruction: it cannot be a metric quantity, a label, an ordering, a part-count, a refinement history, or any piece of local combinatorial structure. This paper proves that what remains has exactly one possible home. If a native obstruction survives at all, it must survive as route-dependence in the way admissible merge and split operations *compose* — and route-dependence of this kind is, by definition, a first-cohomology class.

So the entire remaining question is forced onto a single, concrete object. The merge and split operations, taken together with the relations that say which different descriptions count as "the same move," assemble into a definite mathematical structure. Because every merge can be undone by a split, that structure is a *group*, and the obstruction is precisely the part of that group that survives abelianisation.

That much was the paper's original goal: determine whether admissible refinement motion could carry a surviving route-dependence residue at all. But the closure (Gate-3) programme has since changed what the right question is. It independently identified a *specific* closure-transport residue — a sevenfold holonomy κ carried by a rotation-label cochain on the vacuum transport complex — and argued that primitive Facts generate protected, non-shrinkable loops carrying exactly this winding. This paper therefore takes the stronger view. The question is no longer "does refinement motion manufacture some residue?" but "does refinement motion *detect the same residue* Gate-3 already found?" A residue that must be built from scratch is speculative; a residue already exhibited elsewhere in the programme, which refinement loops need only be shown to register, is concrete.

There is a reason such a loop, once found, could never be smoothed away — and it is the reason the trace would be permanent. When reality commits a fact, it isolates a region and discards it. A loop drawn around that discarded region cannot be shrunk to a point, because shrinking it would sweep across the discarded region and reach back into it — and reaching back in would undo the commitment, the one thing that cannot be undone. So the loop is not held open by some delicate conserved quantity that might quietly switch off; it is held open by the irreversibility of the fact itself. That is what makes the trace unerasable, and it is why the permanence is structural rather than a matter of lasting through time.

What all of this leaves is a single, sharp question — no longer "is there a hidden obstruction somewhere?" but one concrete property of one specific map. Reality's committed loops live on the closure side; refinement motion lives on the reconstruction side; the two are joined by the programme's transport map. The whole remaining uncertainty is whether that map carries one ordinary, reversible refinement loop onto the committed loop around a discarded region. If it does, the trace is real, permanent, and the very object the closure programme found from the other side — and reality keeps its small unerasable tally of genuine commitment. If it does not, every refinement loop smooths to nothing and no such tally survives. The paper proves there is nowhere else for the answer to hide; what it cannot yet do is reach across the gap between a reversible refinement and an irreversible fact, and that one crossing is the work that remains.

*Epistemic markers: **(established)** for results inherited from prior VERSF papers or standard mathematics; **(proven)** for results proved here; **(conditional)** for results holding under a stated assumption; **(conjectural)** / **(open)** for what remains undecided.*

Abstract

The Admissible Lift Theorem established that under FP1, IA, and full admissible merge/split equivalence, all admissible finite refinements of a fixed unresolved region lie in one connected component, so $H^0(\mathcal{Q}\_full) = \mathbb{Z}$ and the component-seal (H^0) branch of Reversible-Connectedness failure is eliminated natively. The present paper addresses the only remaining native residue.

The contribution is structural, not computational. We prove a **localisation theorem**: every surviving native RC residue is a route-dependence (holonomy) phenomenon and therefore factors through the first cohomology of the merge–split composition complex $\mathcal{C}$_MS. With H^0 closed, no single-valued seal can survive; the only surviving mechanism is a quantity conserved along every move yet nontrivial around a closed loop — exactly an H^1 class. There is no remaining native location outside admissible merge–split composition structure.

We then sharpen the object. Because FP1 makes every merge reversible by a split, the merge–split category MS(M) is a **groupoid**; its connectedness (inherited) gives $\mathcal{C}\_MS \simeq BG$ for a single group $G = \pi_1(\mathcal{C}\_MS, \{M\})$, the group of closed merge–split words modulo coherence. Hence

$H_1(\mathcal{C}\_MS) = G\_ab$, $H^1(\mathcal{C}\_MS; A) = \mathrm{Hom}(G\_ab, A)$,

so the **Residue Reduction Theorem** holds: the native RC residue *is* G_ab, and survives iff G is not perfect — most simply closing when G is trivial and $\mathcal{C}$_MS is contractible. A necessary-and-sufficient criterion follows: $H_1(\mathcal{C}\_MS) = 0$ iff every closed merge–split word is a product of coherence relations. Three relation classes are *forced* — cancellation (FP1 reversibility), presentation coherence (IA), and exchange (independence of disjoint regions) — and we prove a Residual Topology Theorem: any class surviving these is genuinely global, not bookkeeping. What is **not** forced, and is the single open question, is higher-order coherence exhaustion: whether the remaining relations collapse G_ab to zero.

Finally, universal coefficients make the Gate-3 contact precise, and reframe it from generation to **detection**. A sevenfold cycle ($w^7 \simeq \mathrm{id}$, $w \not\simeq \mathrm{id}$) would put a $\mathbb{Z}_7$ subgroup in G_ab, *invisible to $H^1(\cdot; \mathbb{Z})$ but visible to $H^1(\cdot; \mathbb{Z}_7)$* — the same torsion-versus-coefficient question the closure programme reached from the other side. The closure programme has independently produced such a class, $\kappa \in H^1(\Gamma\_vac; \mathbb{Z}_7)$, carried by primitive-Fact cycles. The remaining question is therefore not whether refinement motion *creates* a residue but whether its loops *detect* κ. We isolate the comparison map $\Phi : H^1(\mathcal{C}\_MS; A) \to H^1(\Gamma\_vac; \mathbb{Z}_7)$ as the object whose *full* construction the isomorphism question would need, while detection needs only its evaluation on a single loop, and prove a conditional Primitive Fact Detection theorem: if the primitive-Fact loop is realised in refinement motion and carries nonzero κ-charge, then G_ab ≠ 0 — the residue

is real, Coherence Exhaustion fails, and the surviving native class *is* the κ-charge seen from the reconstruction side.

Table of Contents

---

1. Introduction

The companion sequence progressively reduced what first appeared to be a large family of unrelated residual questions. Reversible Connectedness (RC), the Admissible Lift Property (ALP), operational sectors, refinement topology, quotient geometry, admissible transport, and closure topology were shown to be different manifestations of one underlying problem.

The Admissible Lift Theorem completed the previous reduction. Under full admissible merge/split equivalence it established

$H^0(\mathcal{Q}\_full) = \mathbb{Z}$ **(established)**,

so the refinement quotient is connected: any admissible refinement reaches any other. This eliminates the **degree-zero** obstruction — the *component-seal*, in which admissible motion is trapped in a sector it cannot leave. The reachability content of RC therefore holds natively.

But connectedness does not imply contractibility. A connected space may still carry nontrivial loops, and the entire remaining native residue is concentrated in one place:

$H^1(\mathcal{Q}\_full)$, equivalently $H^1(\mathcal{C}\_MS)$.

It is worth separating two readings of "RC" that the H^0 result forces apart, because the whole of this paper lives in the gap between them:

- **RC_reach** — *reachability*: any two admissible refinements are joined by some admissible path. This is the H^0 / π_0 statement, and it is **closed** (Admissible Lift Theorem).
- **RC_path** — *route-independence of the closure-charge lift*: any two admissible paths between the same refinements assign the same closure charge, so the **abelian, closure-charge-valued lift** (the $\mathbb{Z}_7$/phase-valued lift carried by transport, §10) is single-valued. This is an H^1 statement, and it is the residue.

A precision the rest of the paper depends on, stated here once. The lift whose single-valuedness is at issue is the *abelian* one — it takes values in the closure charge group $\mathbb{Z}_7$ (a phase), and the obstruction to its single-valuedness is the monodromy representation $\pi_1(\mathcal{C}\_MS) \to \mathbb{Z}_7$. Because $\mathbb{Z}_7$ is abelian, that representation factors through the abelianisation G_ab, so **G_ab is the complete obstruction for the closure-charge lift** — Hom(G, $\mathbb{Z}_7$) = Hom(G_ab, $\mathbb{Z}_7$), and nothing in the commutator subgroup [G, G] can carry a closure charge. This is why the whole downstream analysis (Theorem 3.5, the H^1 criterion, the "G perfect" closing branch) controls exactly the right object: for the $\mathbb{Z}_7$-lift, G_ab is not a shadow of the obstruction but the obstruction itself. The residual *non-abelian* monodromy — single-valuedness of a hypothetical non-abelian-valued lift, which would need G = 1 rather than merely G_ab = 0 — is a strictly separate question, carrying no closure charge, and is recorded as such in the open set (§12); it is not part of RC_path as used here.

The reconstruction needs the closure-charge lift to be well-defined, which is RC_path in this abelian sense, not merely RC_reach. So the question is no longer "can refinements reach one another?" (yes) but "do different admissible routes between them carry the same closure charge?"

The purpose of this paper is not to ask whether such a residue exists. It is to determine *where it could possibly live*, and the answer is the central theorem: **every surviving native RC_path residue factors through the composition structure of admissible merge and split operations, and through nothing else.** The remainder constructs the relevant complex, proves the localisation theorem, sharpens it via the groupoid structure forced by reversibility, identifies which coherence relations are forced, derives the H^1 criterion, and isolates the single comparison that would identify the residue with the Gate-3 closure class κ.

---

2. The Merge–Split Localisation Theorem

This section proves that any surviving native RC residue must appear as a first-cohomology class of the merge–split composition complex. The statement is stronger than a computational proposal: it is a *localisation* theorem, identifying the unique location a surviving native residue can occupy.

2.1 What has already been eliminated (established)

The companion sequence eliminated, one by one, every candidate native obstruction:

- external labels — excluded by IA;
- capacity-derived invariants — mathematically constant under capacity transfer, hence no seal;
- ordering data — presentation-dependent, removed by admissible relabelling;
- fixed arity (part-count) — survives only as an added dynamical seal (Theorem 9.1B of the companion), not as a native invariant;
- refinement history — quotiented away as presentation;
- local combinatorial structure (valence, adjacency, branching, tree depth) — removed by admissible coarsening;
- disconnected components (the H^0 component-seal) — eliminated under full merge/split admissibility, $H^0(\mathcal{Q}\_full) = \mathbb{Z}$.

Consequently no surviving native obstruction can depend on any *local* refinement datum. Any surviving residue must be **global**, and the only remaining admissible structure is the way admissible refinement operations compose. This is the observation the localisation theorem makes precise.

2.2 From single-valued seals to holonomy — the $H^0 \to H^1$ step (proven)

The decisive step is the passage from "the residue is global" to "the residue is specifically H^1." It runs as follows.

A conserved seal is, by definition, invariant under every admissible move; it is therefore **locally constant** on the quotient $\mathcal{Q}$_full. A locally constant function on a connected space is globally constant. But $\mathcal{Q}$_full *is* connected ($H^0 = \mathbb{Z}$), so every single-valued, locally-constant seal is globally constant — it separates nothing. The component-seal mechanism is dead.

The only way admissible structure can still carry a *nontrivial* invariant is therefore not as a single-valued function but as a **holonomy**: a quantity unchanged along every move, yet whose accumulated value around a closed admissible loop is nontrivial — a route-dependence rather than a labelling. Holonomy of the admissible-move structure is, by definition, a class in $H^1(\mathcal{C}\_MS)$. This is exactly the RC_path obstruction of §1: reachability intact, route-equivalence failing.

We record the one honest caveat, then immediately narrow it. The argument forces a *single-valued* seal into H^0 (now closed) and a *route-dependent* seal into H^1. It does not, by itself, exclude a residue carried by a higher-degree mechanism — an H^{2+} class acting through some more elaborate global structure. That degree-exhaustion question (§7, §12) is inherited unchanged from the companion's §9.8, and as an abstract question about whether $\mathcal{C}$_MS carries seal-bearing higher cohomology it remains open. But it does not touch the object this paper is actually about, and the following proposition says why.

Proposition 2.3 (Degree separation) (proven)

Any invariant detected by transport along loops — in particular the closure-charge lift of §1, valued in $\mathbb{Z}_7$ — factors through $\pi_1(\mathcal{C}\_MS)$, hence through $H_1(\mathcal{C}\_MS)$ = G_ab, and is therefore insensitive to $H^k(\mathcal{C}\_MS)$ for every $k \geq 2$.

Proof. The closure charge is a class $\kappa \in H^1(\Gamma\_vac; \mathbb{Z}_7)$ evaluated on transported loops; as a degree-one cohomology class it pairs with H_1 and vanishes on boundaries, so its value on a loop depends only on the loop's class in $H_1(\mathcal{C}\_MS)$ (equivalently, by homotopy-invariance of holonomy of a flat/coherence-respecting connection, on its class in π_1). A loop-valued pairing is a function of 1-cycles; an H^k class with $k \geq 2$ is a function of k-cycles (surfaces and higher), on which no loop pairing is defined. Hence no H^{2+} class contributes to the closure-charge lift.

The consequence is that the H^{2+} caveat, real as a pure-mathematical question about $\mathcal{C}$_MS, is *irrelevant to the closure-charge lift, to κ-detection, and to RC_path as defined in §1*. Wherever this paper qualifies a closure-charge statement "modulo H^{2+}," that qualification is discharged by Proposition 2.3: the closure-charge residue lives entirely in H^1 = G_ab, and a higher-degree mechanism could only carry some *other* kind of seal, one that by construction carries no closure charge and is invisible to κ. So the localisation, for the lift at issue, is to H^1 exactly — not "H^1 modulo a higher possibility." Within the route-dependence mechanism — the only one that can carry a closure charge — the residue is precisely H^1.

Theorem 2.1 (Merge–Split Localisation) (proven; H^{2+} caveat discharged for the closure-charge lift by Prop 2.3)

Let $\mathcal{Q}$_full be the fully reduced refinement quotient under admissible equivalence, and let $\mathcal{C}$_MS be the merge–split composition complex generated by admissible merge and split operations together with their admitted coherence relations. Then every surviving native RC_path residue is represented by a nontrivial class of

$H^1(\mathcal{C}\_MS)$.

Equivalently, modulo coherence-exhaustion (§7) and the H^{2+} degree question (§2.2), a native RC_path residue survives **iff** $H^1(\mathcal{C}\_MS) \neq 0$.

Proof. By §2.1 the eliminations remove all local admissible data, so any surviving native invariant descends to the connected quotient $\mathcal{Q}$_full. By §2.2 a single-valued seal is then globally constant ($H^0 = \mathbb{Z}$), so a nontrivial survivor cannot be single-valued; it must be a holonomy of the admissible-move structure. The admissible-move structure is generated entirely by merge and split operations modulo their coherence relations, and its holonomies are by construction the first-cohomology classes of the complex $\mathcal{C}$_MS built from those generators and relations. Hence every surviving native residue is a nontrivial class of $H^1(\mathcal{C}\_MS)$. No non-global location remains by §2.1, and no single-valued (H^0) location remains by §2.2.

Corollary 2.2 (proven)

A surviving native RC residue cannot be metric, label-based, ordering-based, arity-based, history-based, presentation-based, or local-combinatorial. If it exists, it is a first-cohomology class of the merge–split composition complex. This completes the localisation programme initiated by the Admissible Lift Theorem: the residue has been cornered from "somewhere in refinement space" to "one cohomology group of one complex."

---

3. The Merge–Split Groupoid and Its Classifying Complex

The localisation theorem says *where* the residue lives; this section sharpens *what kind of object* it is. The key structural fact — forced by FP1 reversibility — is that the merge–split structure is a groupoid, which collapses "compute a cohomology group" to "compute one group and abelianise."

3.1 The category MS(M)

Let M be a fixed unresolved region. Define MS(M):

- **Objects** — admissible refinement classes of M (partitions of M modulo presentation).
- **Generating morphisms** — elementary merges μ (coarsening: $n \to n-1$ parts), elementary splits σ (refining: $n-1 \to n$ parts), and presentation equivalences ρ (relabelling, ordering, history changes).
- **Composition** — ordinary composition of admissible refinement moves.

MS(M) records no metric structure — that has already been quotiented away — only admissible refinement motion and the relations among such motions. This is precisely the structure identified by Theorem 2.1.

3.2 Reversibility makes MS(M) a groupoid (proven, given FP1)

FP1 requires admissible motion to be reversible. Concretely, an elementary split σ that divides one part into two is undone by the elementary merge μ that recombines exactly those two parts, and vice versa:

$\mu \circ \sigma \simeq \mathrm{id}$, $\sigma \circ \mu \simeq \mathrm{id}$ (cancellation, §5.1).

So **every generating morphism is invertible**. Presentation equivalences ρ are invertible by construction. A category in which every morphism is invertible is a groupoid; hence MS(M) is a groupoid.

This is not a modelling convenience but a consequence of the founding reversibility axiom: the irreversibility that produces committed facts lies *downstream* of refinement, at commitment; pre-commitment refinement motion is reversible, and a reversible move-structure is a groupoid.

3.3 Connectedness gives $\mathcal{C}$_MS ≃ BG (proven, given the inherited H^0 result)

A groupoid decomposes into connected components, one per isomorphism class of objects, and each connected component is equivalent to the one-object groupoid of a single automorphism group. The Admissible Lift Theorem established that $\mathcal{Q}$_full is connected — every refinement merges to the coarsest partition {M} and re-splits to any other — so MS(M) has a **single** isomorphism class. Therefore

$\mathrm{MS}(M) \simeq G$, $G := \mathrm{Aut}_{\mathrm{MS}(M)}(\{M\})$,

the group of merge–split loops based at {M} modulo coherence. Passing to classifying spaces, the merge–split composition complex is

$\mathcal{C}_{MS} = B\cdot\mathrm{MS}(M) \simeq BG$ **(proven)**.

The choice of basepoint {M} is canonical: it is the unique coarsest refinement, the terminal-up-to-coherence object reached by exhausting merges.

3.4 Consequence: the residue is the abelianised fundamental group (established, standard)

For the classifying space of a group, homology and cohomology are group (co)homology, and in degree one:

$H_1(\mathcal{C}_{MS}) = H_1(BG) = G_{ab}$ (the abelianisation of G), $H^1(\mathcal{C}_{MS}; A) = \mathrm{Hom}(H_1(\mathcal{C}_{MS}), A) \oplus \mathrm{Ext}(H_0(\mathcal{C}_{MS}), A) = \mathrm{Hom}(G_{ab}, A)$,

the Ext term vanishing because $H_0(\mathcal{C}_{MS}) = \mathbb{Z}$ is free (connectedness). Therefore the **entire native RC residue is G_ab**, and:

$H^1(\mathcal{C}_{MS}; A) = 0$ for all coefficient groups $A \Leftrightarrow G_{ab} = 0 \Leftrightarrow G$ is perfect.

The simplest closing case is $G = 1$ ($\mathcal{C}$_MS simply connected, in fact contractible since BG is a K(G,1)): then RC_path closes outright with no residue. The residue branch is $G_{ab} \neq 0$. This is the precise content behind the companion's "connected but possibly not contractible" remark (§9.11): connectedness fixes H_0; the loop content G_ab is the separate question.

3.5 The Residue Reduction Theorem (proven; H^{2+} caveat discharged for the closure-charge lift by Prop 2.3)

The preceding subsections together collapse the entire localisation programme to a single identification, which deserves to stand as the paper's headline result.

Theorem 3.5 (Residue Reduction)

Under FP1 (reversibility), IA (the elimination programme), and connectedness of $\mathcal{Q}$_full, the native RC residue **is** the abelianised merge–split holonomy group:

Native RC residue = G_ab, where **G = Aut_{MS(M)}({M})** ,

and consequently

Native RC residue survives $\Leftrightarrow$ $G_{ab} \neq 0$ $\Leftrightarrow$ G is not perfect.

Proof. By Theorem 2.1 every surviving native residue is a route-dependence (holonomy) class of $H^1(\mathcal{C}_{MS})$. By §3.2, reversibility makes MS(M) a groupoid; by §3.3, connectedness makes it the one-object groupoid of the single group G = Aut_{MS(M)}({M}), so $\mathcal{C}_{MS} \simeq BG$. By §3.4, $H_1(\mathcal{C}_{MS}) = G_{ab}$ and $H^1(\mathcal{C}_{MS}; A) = \mathrm{Hom}(G_{ab}, A)$. The residue is therefore identified with G_ab, and exists nontrivially iff $G_{ab} \neq 0$, i.e. iff G is not perfect.

What is fixed, and what is relocated

The theorem makes two kinds of openness, previously entangled, cleanly separate.

- **Fixed here (the identification).** *That* the residue equals G_ab is established. This converts every prior question about hidden invariants "somewhere in refinement space" into a question about one named group. There is no closure-charge residue outside G_ab — and by Proposition 2.3 a higher-degree (H^{2+}) mechanism cannot carry a closure charge at all, so the identification is exact for the lift at issue, not merely "modulo H^{2+}."
- **Relocated, not closed (the computation).** *Which* group G_ab is depends on the coherence relations defining G, i.e. on Coherence Exhaustion (§7). Settling those relations is computing the identified object; it is not re-opening its identity. The headline therefore does exactly what the programme set out to do — replace a conceptual doubt ("does an obstruction exist?") with a definite computation ("is this named group perfect?").

Corollary 3.5A (Residue Uniqueness) (proven, modulo coherence presentation)

There is at most one native closure-charge residue, and it is the single canonical group G_ab. Precisely: every native route-dependence residue is a class of $H_1(\mathcal{C}\_MS)$ = G_ab (Theorem 2.1), and G_ab is canonical — it is the abelianisation of G = Aut_{MS(M)}({M}), well-defined on the nose because the connected-groupoid basepoint {M} is unique up to isomorphism and abelianisation kills the resulting inner automorphisms. Hence no second, independent native residue can exist: any obstruction a later analysis might uncover is *already* an element of G_ab, not a rival to it. In plain terms — **if a residue survives, there is only one.**

This forecloses a specific worry: that future work might turn up a hidden second residue disjoint from the merge–split holonomy. It cannot, within the relevant degree — every closure-charge obstruction is accounted for by G_ab, and by Proposition 2.3 nothing in higher degree carries a closure charge to compete. (The bare-mathematical question of whether $\mathcal{C}$_MS has seal-bearing H^{2+} of some *other*, charge-free kind is separate and stays in the open set, §12; it is not a second closure-charge residue.)

Everything after this section is, in effect, the analysis of G_ab: §§4–5 present its generators and forced relations, §6 gives the perfect/non-perfect criterion, §7 isolates the one open input, and §8 fixes the signature (order-7 torsion) under which G_ab meets the Gate-3 class κ.

3.6 Detection versus Generation

The Residue Reduction Theorem identifies the native RC residue with G_ab, the abelianisation of the merge–split holonomy group. Read on its own, this poses a *generation* question: does admissible refinement motion, by itself, manufacture a nontrivial holonomy?

The Gate-3 programme supplies a reason to ask a sharper question. It produces, independently of refinement motion, a holonomy object on the vacuum transport complex:

$\kappa = [\rho\_rot] \in H^1(\Gamma\_vac; \mathbb{Z}_7)$ **(established in the Gate-3 papers)**,

where ρ_rot is the closure-transport rotation-label cochain. *(Notation: the Gate-3 papers write this cochain ρ; we write ρ_rot here to avoid collision with the presentation-equivalence morphism ρ of §3.1.)* The later Gate-3 results argue further that primitive Facts generate *protected* non-bounding cycles carrying nonzero closure winding — so κ is not merely a formal class but one with concrete charged representatives. We take this as an inherited Gate-3 input, carrying whatever marker it holds in its home paper; the detection and persistence results below are conditional on it accordingly.

This reframes the comparison the rest of the paper pursues. The merge–split programme need not **generate** κ. It need only **detect** it.

- *Generation question (weaker target, now superseded):* "Does admissible refinement motion create a sevenfold residue?" — speculative, because it asks refinement motion to produce structure from nothing.

- *Detection question (the right target):* "Do admissible refinement loops carry nontrivial κ-charge?" — concrete, because the charge already exists; the only question is whether refinement loops register it.

The distinction matters because the two have different burdens of proof. Generation would require exhibiting a nontrivial G_ab purely from the coherence combinatorics (the Coherence-Exhaustion computation of §7). Detection requires only a *pairing*: a way to evaluate the already-existing class κ on refinement loops, and a single loop on which it is nonzero. As Corollary 10.6 makes precise, any such nonzero pairing factors through G_ab (because its target $\mathbb{Z}_7$ is abelian) — so detection, if it holds, settles the generation question as a corollary, in the affirmative. We develop the detection route in §10; the present subsection only fixes that detection, not generation, is the question Gate-3 has made available.

---

4. Closed Words and Holonomy

The localisation theorem identifies the residue with $H^1(\mathcal{C}\_MS)$ = Hom(G_ab, A); §3 identifies G_ab as the abelianised group of merge–split loops. This section names the generators of that group — closed words — and fixes the holonomy interpretation.

4.1 Closed words

A **merge–split word** is a finite composition of admissible elementary moves,

$w = m_1\, m_2 \cdots m_n$, each $m_k \in \{\mu, \sigma, \rho\}$.

A word is **closed** if it begins and ends at the same refinement class,

$R \to R_1 \to R_2 \to \cdots \to R$,

so a closed word is a loop in $\mathcal{C}$_MS, i.e. an element of $\pi_1(\mathcal{C}\_MS, R) = G$ (for $R \simeq \{M\}$). The central question becomes: *can every closed word be reduced to the identity by admissible coherence relations?* If yes, G = 1 and no residue survives; if not, a surviving loop represents a candidate H^1 class.

4.2 Holonomy interpretation — reachability versus route-equivalence (proven)

Closed words measure route-dependence. Let two admissible paths join the same refinements,

$P_1 : R \to R'$, $P_2 : R \to R'$,

then $P_1 P_2^{-1}$ is a closed word (the inverse exists by §3.2). If $P_1 P_2^{-1}$ reduces to the identity, the two routes are equivalent; if not, a residual **holonomy** survives. This residual holonomy is exactly the obstruction RC_path of §1:

- **Reachability** (RC_reach, H^0) remains intact — both P_1 and P_2 exist.
- **Route-equivalence** (RC_path, H^1) is what can fail — P_1 and P_2 need not agree.

The distinction is the whole subject of the paper. The residue does not concern whether refinements can reach one another; it concerns whether the admissible lift along different routes is single-valued.

4.3 Closed words as (co)homology generators, and universal coefficients (established)

Closed words generate the cycle space of $\mathcal{C}$_MS; the coherence relations (2-cells) generate the boundaries. First homology is the quotient:

$H_1(\mathcal{C}\_MS) = Z_1 / B_1 = \langle\text{closed words}\rangle / \langle\text{coherence relations}\rangle = G\_ab$,

the surviving closed merge–split words after admissible reduction. Dually, by universal coefficients,

$H^1(\mathcal{C}\_MS; A) = \mathrm{Hom}(G\_ab, A)$.

The coefficient group A is not a formality here, and this is the point that connects to Gate-3. If G_ab is *torsion-free*, its torsion-sensitive content is invisible and integral cohomology suffices. But if G_ab carries **torsion** — an element of finite order — that torsion is detected only by cohomology with matching coefficients:

$G\_ab \supseteq \mathbb{Z}_7 \Rightarrow H^1(\mathcal{C}\_MS; \mathbb{Z}) = \mathrm{Hom}(\mathbb{Z}_7, \mathbb{Z}) = 0$, yet $H^1(\mathcal{C}\_MS; \mathbb{Z}_7) = \mathrm{Hom}(\mathbb{Z}_7, \mathbb{Z}_7) \neq 0$.

So a residue can be *cohomologically invisible over $\mathbb{Z}$ and visible over $\mathbb{Z}_7$* — exactly the torsion-versus-coefficient ambiguity the closure programme met from the opposite direction (companion §9.4). We return to this in §8.

---

5. Coherence Relations and Which Are Forced

§4 identifies the objects capable of *generating* a residue (closed words). This section identifies the relations capable of *destroying* one (2-cells), and separates the relations that are **forced** by FP1/IA from those that are conjectural.

The guiding principle: a loop must not generate cohomology merely because the same admissible transformation was described twice. Only genuinely irreducible loops should survive. Each coherence relation contributes a 2-cell; the question of *which* relations are mandatory is what makes the H^1 computation well-posed.

5.1 The relation classes

- **Cancellation.** A split followed by the matching merge is trivial, and conversely: $\mu \circ \sigma \simeq \mathrm{id}$, $\sigma \circ \mu \simeq \mathrm{id}$ (when the original subdivision is restored). Removes backtracking.
- **Associativity (merge).** Merging three regions in different orders agrees: $(P_1 \cup P_2) \cup P_3 \simeq P_1 \cup (P_2 \cup P_3)$. A 2-cell; any loop built only from re-association bounds.
- **Coassociativity (split).** The dual statement for nested splits; distinct nested splitting sequences with the same final refinement are identified.
- **Exchange.** Moves on disjoint refinement sectors commute: $m_1 m_2 \simeq m_2 m_1$ for independent m_1, m_2.
- **Presentation coherence.** Relabelling, ordering changes, and history changes carry no content; loops built solely from them are declared trivial.

5.2 Which relations are forced (proven)

Three of these are not optional design choices; they are forced by the founding structure, and this is what gives the Residual Topology Theorem (§6) its teeth.

- **Cancellation is forced by FP1.** FP1 makes pre-commitment refinement motion reversible; cancellation is precisely the statement that a split's inverse is the matching merge (§3.2). Denying it would deny reversibility. **Forced.**
- **Presentation coherence is forced by IA.** IA individuates refinements internally and supplies no external indexing; relabelling, ordering, and history are by IA non-physical, so loops built from them must bound. **Forced.**
- **Exchange is forced by independence of disjoint regions.** Operations on disjoint unresolved sub-regions act on independent data; their order cannot carry content, on pain of manufacturing route-dependence from nothing. **Forced.**

Associativity and coassociativity are likewise forced as the on-the-nose associativity of region union and its dual — set union is associative, so merge-association is a genuine identity and contributes a forced 2-cell. What is **not** forced by any of the above is whether these relations *exhaust* the coherences — whether every commuting diagram of mixed merge/split compositions is generated by them, or whether additional higher compatibilities are required. That is the single open question (§7).

5.3 The coherence principle

The forced relations share one interpretation: they identify different *descriptions* of the same admissible transformation. The resulting 2-cells are not added arbitrarily — they are exactly the relations needed to prevent presentation artefacts from masquerading as topology. Consequently any surviving loop, having survived all the forced relations, must represent genuine global structure rather than bookkeeping. This is made precise in §6.

A note on direction, since it controls the honesty of every claim below. Adding relations (2-cells) *enlarges* the boundary subgroup B_1 and therefore *shrinks* $H_1 = Z_1/B_1$. Computing with only the *known forced* relations therefore yields an **upper bound** on G_ab: the true residue can only be smaller as further coherences are found. A class surviving the forced relations is a genuine candidate; whether it survives the full (possibly larger) relation set is the open Coherence-Exhaustion question. We never claim the computed H_1 is exact until exhaustion is settled.

---

6. The H^1 Criterion and Its Rigidity Consequences

Theorem 6.1 (Merge–Split H^1 Criterion) (proven)

Let $\mathcal{C}$_MS be the merge–split composition complex with 1-cells the elementary moves and 2-cells the admissible coherence relations. Then

$H_1(\mathcal{C}\_MS) = 0 \Leftrightarrow$ every closed merge–split word lies in the subgroup generated by the coherence relations,

and consequently, via universal coefficients (§4.3),

$H^1(\mathcal{C}\_MS; A) = 0$ for all A $\Leftrightarrow$ G_ab = 0 $\Leftrightarrow$ G is perfect.

Proof. Closed words generate the cycle space Z_1; coherence relations generate the boundaries $B_1 = \mathrm{im}(\partial_2)$. A closed word is null-homologous exactly when it lies in B_1, so all closed words bound iff $Z_1 = B_1$ iff $H_1 = Z_1/B_1 = 0$. By §3.4, $H_1(\mathcal{C}\_MS)$ = G_ab, so this is G_ab = 0, i.e. G perfect; and $H^1(\mathcal{C}\_MS; A)$ = Hom(G_ab, A) vanishes for all A iff G_ab = 0.

We state the criterion in homology first, deliberately. The integral cohomology $H^1(\mathcal{C}\_MS; \mathbb{Z})$ = Hom(G_ab, $\mathbb{Z}$) sees only the free rank of G_ab and is **blind to torsion**; the homological/coefficient-sensitive statement is the correct one, and is what the sevenfold question (§8) requires.

Corollary 6.2 (proven)

A surviving native RC residue exists iff at least one closed merge–split word survives admissible reduction — equivalently iff G_ab $\neq$ 0. The localisation theorem is thereby operational: the remaining RC question is a finite presentation problem for one group.

Theorem 6.3 (Residual Topology Theorem) (proven, given the forced relations of §5.2)

Assume the forced relations — cancellation (FP1), presentation coherence (IA), exchange (independence). Then any surviving nontrivial class in $H_1(\mathcal{C}\_MS)$ arises from a genuinely global composition cycle; it cannot arise from local refinement bookkeeping.

Proof. Cancellation kills immediate-backtracking cycles. Presentation coherence kills relabelling, ordering, and history cycles. Exchange kills cycles generated solely by independent commuting operations. A surviving cycle is therefore generated by none of: local reversal, local presentation change, local independence. What remains is a global obstruction to contraction — a cycle that does not bound for topological rather than representational reasons.

This sharpens the localisation considerably: the residue is not merely in $\mathcal{C}$_MS, it is in the *genuinely topological sector* of $\mathcal{C}$_MS. $H_1(\mathcal{C}\_MS)$ cannot survive because of bookkeeping; it can survive only because $\mathcal{C}$_MS — equivalently the group G — has genuine global structure.

Theorem 6.4 (Rigidity) (proven, modulo §2.2 degree caveat)

Assume FP1, IA, and the eliminations of the Admissible Lift Theorem. Then every surviving native RC residue is represented by a nontrivial class of $H^1(\mathcal{C}\_MS)$, and no alternative native residue exists.

Proof. By Theorem 2.1 every surviving residue factors through $\mathcal{C}$_MS; by Theorem 6.1 every nontrivial residue corresponds to a surviving closed word, i.e. a nontrivial element of G_ab; by §2.1 all other candidate locations are eliminated. Hence the residue, if any, is a nontrivial H^1 class and nothing else (subject to the H^{2+} degree question of §2.2).

This is the structural heart of the paper. Its purpose is not to compute the final residue but to prove the residue has **nowhere else to hide**. Before it, one could still imagine unknown admissible invariants living somewhere in refinement space; after it, the only admissible place left is G_ab, the abelianised group of merge–split holonomies.

---

7. The Higher-Coherence Problem

Theorems 6.3–6.4 leave a single uncertainty, and it is sharp. It is not cancellation, exchange, or presentation — those are forced and their cycles bound. It is whether the *remaining* coherences are complete: whether every admissible composition diagram commutes once the forced relations are imposed, or whether some diagram fails to commute and thereby leaves a cycle that does not bound.

In the groupoid language of §3, this is whether the relations of §5 give a *complete presentation* of G. If they do, G (hence G_ab, hence the residue) is determined; if additional forced coherences exist, the true G is a further quotient and G_ab can only shrink (§5.3).

Definition 7.1

A **coherence cycle** is a closed merge–split word generated entirely by different admissible compositions of the same overall refinement transformation. Representative families: nested merge structures, nested split structures, mixed merge/split compositions, and repeated refinement/coarsening ladders. These name the *redescription* loops the forced relations of §5.2 do not already remove; they are illustrative families, not an exhaustive description of every closed word. In particular a genuinely topological closed word — γ_D, say — is *not* a coherence cycle in this sense, which is exactly why the conjecture below must quantify over all closed words, not only over coherence cycles.

Conjecture 7.2 (Coherence Exhaustion) (conjectural)

Every closed merge–split word lies in the subgroup generated by the coherence relations — cancellation, associativity, coassociativity, exchange, and presentation coherence — equivalently, every cycle bounds. (This is the all-closed-words quantifier of Theorem 6.1, not a statement about coherence cycles alone; a surviving non-redescription cycle such as γ_D would falsify it.) If true,

$G_{ab} = 0$, hence $H^1(\mathcal{C}_{MS}; A) = 0$ for all A, and RC_path closes (the closure-charge lift is single-valued).

If false, a residual class survives in G_ab and a holonomy residue exists. Because the quantifier is over all closed words, Conjecture 7.2 holds **iff** $G_{ab} = 0$; this is the biconditional on which the mutual exclusivity of Outcomes A and B (§10.4, §13) rests, and it would fail under the narrower coherence-cycle-only reading, where a topological survivor could coexist with all redescription cycles bounding. The entire remaining RC question reduces to the truth or falsity of this single conjecture.

The honest status is that the forced relations give an upper bound on G_ab (§5.3); Coherence Exhaustion is the claim that no further relations cut it down past the forced ones — equivalently, that the only way a cycle can survive is genuine topology, never an unrecognised description-identity. The Residual Topology Theorem guarantees survivors are topological *if* they survive; Coherence Exhaustion is the converse-flavoured claim that the recognised relations already capture every non-topological identity. The two together would make the computed H_1 exact.

---

8. The Sevenfold Question

The Gate-3 programme introduces a specific possibility for the surviving cycle: not arbitrary structure, but a cycle of **finite order**, and the privileged candidate order is seven.

Definition 8.1

A closed word w is a **sevenfold cycle** if

$w^7 \simeq \mathrm{id}$ while $w \not\simeq \mathrm{id}$.

Such a cycle is an order-7 element of G, hence of G_ab, generating a $\mathbb{Z}_7$ subgroup. (We say *subgroup*, not summand: an order-7 element splits off as a direct summand only under extra hypotheses such as finite generation — $\mathbb{Z}_7 \subset \mathbb{Z}_{49}$ is not a summand — and the detection statement below is arranged not to need splitting.)

Proposition 8.2 (proven)

Suppose a sevenfold cycle is *detected*: a closed word w on which the closure-charge homomorphism takes an order-7 value, i.e. the $\mathbb{Z}_7$-valued charge of [w] is nonzero (the situation Corollary 10.6 produces under refinement realisation). Then that charge is itself an explicit nonzero homomorphism G_ab → $\mathbb{Z}_7$, so

$H^1(\mathcal{C}\_MS; \mathbb{Z}_7) = \mathrm{Hom}(G\_ab, \mathbb{Z}_7) \neq 0$,

while $H^1(\mathcal{C}\_MS; \mathbb{Z}) = \mathrm{Hom}(G\_ab, \mathbb{Z})$ may vanish if G_ab is torsion (since Hom of finite-order content into $\mathbb{Z}$ is zero).

Proof. The closure charge $\kappa\_\tau = \kappa \circ \tau\_*$ is a homomorphism G_ab → $\mathbb{Z}_7$ (it factors through the abelianisation because $\mathbb{Z}_7$ is abelian). If it is nonzero on [w], it is a nonzero element of $\mathrm{Hom}(G\_ab, \mathbb{Z}_7) = H^1(\mathcal{C}\_MS; \mathbb{Z}_7)$. Taking A = $\mathbb{Z}$ instead kills all torsion, since $\mathrm{Hom}(T, \mathbb{Z}) = 0$ for any torsion group T.

We lean on the *detected* charge rather than on a bare order-7 element of G_ab, because the bare-element route needs more. An order-7 element of an arbitrary abelian group need not admit any nonzero homomorphism to $\mathbb{Z}_7$ — the Prüfer group $\mathbb{Z}(7^\infty)$ has elements of order 7 yet $\mathrm{Hom}(\mathbb{Z}(7^\infty), \mathbb{Z}_7) = 0$. The conclusion from a bare order-7 element therefore requires G_ab finitely generated; the detection route avoids the hypothesis entirely by handing us the homomorphism directly.

Significance

This is the first point where the RC programme and Gate-3 can meet *algebraically* rather than heuristically. The comparison becomes: **does $\mathcal{C}$_MS contain a sevenfold composition cycle?** If yes, a Gate-3-compatible residue exists — a $\mathbb{Z}_7$ holonomy detected exactly in $\mathbb{Z}_7$ coefficients. If not, the Gate-3 residue must arise elsewhere. Two senses of "sevenfold cycle" are in play and should be kept distinct: the bare group-theoretic one (Definition 8.1 — an order-7 element of

G_ab exists) fixes the *algebraic signature* of a candidate residue, while the detection sense of Proposition 8.2 (the closure charge is nonzero on [w]) is what converts that signature into *visible $\mathbb{Z}_7$ cohomology*. They can come apart — a bare order-7 element could carry zero κ-charge — which is why Proposition 8.2 conditions on the stronger, detected form; the bare element alone fixes the shape, detection supplies the witness.

Proposition 8.2 also explains *why* the integral and mod-7 statements can disagree, which is not a technicality but the crux of the cross-programme comparison. The companion's §9.4 reached the same wall from the closure side: whether the Gate-3 $\mathbb{Z}_7$ is genuine homological torsion in H_1(Γ_vac; $\mathbb{Z}$) or merely a coefficient choice. Here the identical question is whether the surviving merge–split holonomy is genuine torsion in G_ab or absent over $\mathbb{Z}$. The two programmes' open questions are not merely the same *shape*; under the groupoid identification they are the same *kind of torsion-detection question about a first homology group*. §10 turns this from a comparison of two classes into a detection test: whether refinement loops register the κ-charge Gate-3 has already exhibited.

Nothing here proves a sevenfold cycle exists — it is **(open)** whether one does. The proposition only fixes its signature: order-7 element of G_ab, visible in $\mathbb{Z}_7$ cohomology, invisible in $\mathbb{Z}$ cohomology.

---

9. Toy Models

These models are not proofs; they illustrate the two branches the criterion of §6 distinguishes.

9.1 Complete coherence — the closing branch

Suppose every closed merge–split word reduces, modulo the coherence relations, to a product of cancellation, associativity, coassociativity, exchange, and presentation coherence (Conjecture 7.2 holds). Then every cycle bounds, G_ab = 0, and

$H^1(\mathcal{C}$_MS; A) = 0 for all A.

No admissible holonomy survives and RC_path closes. This is the closing branch in its general form, G_ab = 0 (G perfect); the *detectable* residue is gone and κ is undetected. The stronger picture — $\mathcal{C}$_MS contractible, the refinement quotient fully coherent in every degree — is the sub-case G = 1, in which the higher-degree (H^{2+}) question closes as well; the merely-perfect case leaves that question open (a perfect nontrivial G has nonzero Schur multiplier H_2). **RC_path closes outright; full contractibility is the G = 1 specialisation.**

9.2 Residual holonomy — the residue branch

Suppose there is a closed word w with $w \not\simeq id$ but $w^n \simeq id$ for some finite n. Then w is a nontrivial torsion element of G_ab; the cycle survives admissible reduction, and

$H^1(\mathcal{C}\_MS; \mathbb{Z}_n) \neq 0$.

A holonomy residue survives; the refinement quotient retains route-dependence. **RC_path fails, with a finite-order seal.**

9.3 The sevenfold residue — the κ-compatible branch

The case n = 7 is distinguished: $w^7 \simeq id$, $w \not\simeq id$, generating a $\mathbb{Z}_7$ in G_ab. The residue is then directly compatible with the Gate-3 closure sector, detected in $H^1(\mathcal{C}\_MS; \mathbb{Z}_7)$ and invisible in $H^1(\mathcal{C}\_MS; \mathbb{Z})$ (Proposition 8.2). Nothing here proves such a cycle exists; the model exhibits the precise structure a κ-compatible branch would require.

---

10. Relation to Gate-3

The localisation fixes the RC residue in $H^1(\mathcal{C}\_MS)$. The Gate-3 programme independently identifies a closure residue $\kappa \in H^1(\Gamma\_vac; \mathbb{Z}_7)$. This section states the relationship as a single, well-posed comparison rather than a conjectured identity.

10.1 The two residues

	RC route	Gate-3 route
Origin	admissible refinement motion	closure topology and transport
Residue	$H^1(\mathcal{C}\_MS)$; = Hom(G_ab, A)	$\kappa \in H^1(\Gamma\_vac; \mathbb{Z}_7)$
Mechanism	merge–split holonomy	closure holonomy

The two arise from different starting points, yet both terminate in **first cohomology** — and, under §3, both are first homology / abelianised-holonomy questions. That is what makes the comparison precise rather than analogical.

10.2 The comparison problem, recast as detection (open, with a conditional result)

The neutral form of the comparison asks whether there is a natural map

$\Phi : H^1(\mathcal{C}\_MS; A) \to H^1(\Gamma\_vac; \mathbb{Z}_7)$

sending the RC residue to κ — equivalently a homomorphism $G\_ab \to H_1(\Gamma\_vac; \mathbb{Z})$ carrying the merge–split holonomy onto the closure torsion. That remains the object to be constructed. But §3.6 sharpened the question from generation to detection, and detection is the easier and more concrete half of Φ: it does not require the full isomorphism, only that κ can be *evaluated on* refinement loops and is nonzero on at least one.

Make this precise. To pair the closure class $\kappa \in H^1(\Gamma\_vac; \mathbb{Z}_7)$ with a merge–split loop, one needs a map carrying refinement cycles into Γ_vac. That map is not a new object invented for the comparison: it is the programme's **admissible transport**, the same apparatus through which the closure cochain ρ_rot is defined (§3.6). Write τ for admissible transport on refinement words. The **κ-charge** of a closed merge–split word w is then

$\kappa\_\tau(w) := \kappa(\tau\_* w) \in \mathbb{Z}_7$,

the closure class evaluated on the transport image of the loop. The open content of the full bridge Φ is precisely *which* closed words admit a transport image and what it is; detection asks only the much smaller question: is there one closed word w with $\kappa\_\tau(w) \neq 0$?

10.2A Properties required of admissible transport

The endgame leans on τ, so it is worth stating exactly which properties of admissible transport the theorems below use — not deriving them here, but making the interface to the inherited machinery explicit, so that what is assumed of τ is visible rather than tacit. Every result from here uses only these three.

1. **Functoriality.** τ respects composition of admissible moves: $\tau(f \circ g) = \tau(f) \circ \tau(g)$. (This is what makes τ a functor MS(M) → (the closure-transport category), and hence what lets it act on words at all.)
2. **Boundary compatibility (chain map).** τ commutes with the boundary operator: $\partial \circ \tau\_* = \tau\_* \circ \partial$. This is the property the persistence proof (Theorem 10.8) uses to move a filling of a refinement loop in $\mathcal{C}\_MS$ to a filling of its image in Γ_vac; it follows from functoriality once τ is realised on cells.
3. **Closure-charge preservation.** τ carries closure state faithfully, so the charge read on the image equals the charge read on the loop: $\kappa(\tau\_* w) = \kappa\_\tau(w)$. This is not a new assumption — it *is* the faithful-recording premise of Theorems 10.5 and 10.8, named here as a property of τ. Where those theorems assume faithful recording, they assume exactly (3).

With (1)–(3) isolated, the dependence of the endgame on inherited transport is fully exposed: every downstream theorem is conditional on τ having these three properties and on nothing else about it. Functoriality and boundary compatibility are structural (they hold of any transport worthy of the name); charge preservation is the one with physical content, and it is precisely the faithful-recording hypothesis.

10.3 Primitive Fact detection

The Gate-3 programme supplies a candidate charged loop: the **primitive-Fact cycle γ_D**, the loop associated with a single primitive commitment, which the later closure papers argue is a protected non-bounding cycle of nonzero closure winding. The detection question is whether γ_D is reached by transport *from a refinement loop* — and that condition deserves a name, because it is where the entire remaining uncertainty now lives.

Definition 10.4A (Refinement realisation)

A primitive-Fact cycle γ_D is **refinement-realised** if the admissible refinement dynamics contain a closed merge–split word w whose admissible-transport image is γ_D:

$\tau\_* w = \gamma\_D$, w a closed word in $\mathcal{C}\_MS$.

Refinement realisation is thus a single, concrete structural property: that the commitment loop γ_D lies in the transport image of some reversible refinement loop. It is not the assertion that a bespoke bridge exists; the map τ is fixed by the programme, and the only question is whether γ_D is in its image-from-a-loop.

Theorem 10.5 (Primitive Fact Detection) (conditional on refinement realisation)

Assume:

1. *(existence)* primitive Facts exist, and γ_D is the associated closure cycle;
2. *(refinement realisation, Def 10.4A)* γ_D is refinement-realised by a closed merge–split word w, i.e. $\tau\_* w = \gamma\_D$;
3. *(faithful recording)* admissible transport preserves γ_D's closure winding, so $\kappa(\gamma\_D)$ is its genuine nonzero Gate-3 value.

Then the realising word w is a κ-charged refinement loop:

$\kappa\_\tau(w) = \kappa(\tau\_* w) = \kappa(\gamma\_D) \neq 0$,

and consequently admissible refinement motion contains a κ-charged loop.

Proof. By (2) w is a closed word of $\mathcal{C}\_MS$ with $\tau\_* w = \gamma\_D$. By (1) and (3) κ evaluates on γ_D to the nonzero closure winding established in the Gate-3 papers (primitive Facts generate non-bounding cycles of nonzero winding). Hence $\kappa\_\tau(w) = \kappa(\gamma\_D) \neq 0$.

The single hypothesis carrying the uncertainty is now isolated and named: **refinement realisation** (2). Hypotheses (1) and (3) are inherited Gate-3 inputs — that primitive Facts exist and that transport records their winding faithfully. The opaque "a bridge exists on γ_D" is thereby replaced by the definite question of whether a reversible refinement loop transports onto the commitment cycle γ_D. This is exactly the category tension the framework should expect to be the crux: γ_D is a *commitment* object (irreversible, downstream of refinement), while w is a *pre-commitment* reversible loop, and refinement realisation asks whether admissible transport bridges the two. The recast does not close that question — it states it as one concrete property of one named map, which is the most it can honestly be reduced to here.

Corollary 10.6 (Detection forces the residue branch) (conditional)

Under the hypotheses of Theorem 10.5, $G\_ab \neq 0$; equivalently $H_1(\mathcal{C}\_MS) \neq 0$ and Coherence Exhaustion (Conjecture 7.2) is **false**.

Proof. $\kappa\_\tau(\cdot) = \kappa \circ \tau\_*$ is a homomorphism $H_1(\mathcal{C}\_MS) \to \mathbb{Z}_7$ (defined on words admitting a transport image). Since $\mathbb{Z}_7$ is abelian, it factors through the abelianisation, giving a homomorphism $G\_ab \to \mathbb{Z}_7$ that is nonzero on the class of the realising word w. A nonzero homomorphism out of G_ab forces $G\_ab \neq 0$, so $H_1(\mathcal{C}\_MS) \neq 0$. By Theorem 6.1 that is exactly the failure of G to be perfect, i.e. the negation of Conjecture 7.2.

This is stronger than the weaker reading "the image intersects the κ-sensitive sector," and it is worth stating the strength plainly. Detection does not merely *touch* the κ-sector — it **decides** the paper's central open question in the residue direction. If γ_D is refinement-realised and carries its Gate-3 charge, then refinement motion is *not* coherent: a route-dependence residue survives, RC_path does **not** close natively, and the surviving native class carries the κ-charge. The price is honest and named: refinement realisation (Def 10.4A), the question of whether a reversible refinement loop transports onto γ_D. Detection therefore converts "construct all of Φ" into the lighter "exhibit one refinement loop that transports onto γ_D" — a single concrete property of the transport map, not an open-ended search for a bridge.

What Theorem 10.5 does **not** give is the full identification $G\_ab \cong H_1(\Gamma\_vac; \mathbb{Z}_7)$. It establishes a nonzero homomorphism $G\_ab \to \mathbb{Z}_7$ through the κ-charge, not an isomorphism of the two holonomy groups. The residue is shown to be at least as large as the κ-charge it detects; whether it is exactly the closure holonomy — whether merge–split holonomy and closure-transport holonomy are two presentations of one group — is the remaining content of Φ.

Theorem 10.7 (Detection Dominance) (conditional)

If the hypotheses of Theorem 10.5 hold, then the native RC residue is established **independently of Coherence Exhaustion** (Conjecture 7.2): its existence follows from detection alone, without computing the coherence presentation of G.

Proof. Theorem 10.5 gives a nonzero κ-charge $\kappa\_\tau(w) \neq 0$ on the realising word. Corollary 10.6 converts this to $G\_ab \neq 0$, which is the residue's existence and, simultaneously, the negation of Conjecture 7.2. No step invokes Conjecture 7.2 as a premise; the residue is obtained from the detected charge alone.

The force of this is a change of *target*, not a claim of strict superiority. The coherence route (§7) must establish a statement about *all* admissible composition cycles — that the forced and higher relations exhaust them. The detection route needs *one* witness: a single closed refinement word transporting onto γ_D with its Gate-3 charge. A single realising loop is a far smaller object to exhibit than a complete coherence presentation, which is the sense in which detection dominates.

The cost is refinement realisation (Def 10.4A): detection trades the large combinatorial target for the single concrete question of whether the transport map carries some reversible loop onto the commitment cycle γ_D. So the two routes are not strictly ranked — detection is the smaller target but not the free one — and which is the better bet depends on whether refinement-realising γ_D is easier than settling exhaustion. What is unconditional is that *either* route, if completed, ends the question.

Corollary 10.7A (Abelian visibility of a charged loop) (proven)

A detected κ-charged loop is not merely nontrivial in G; it is *visible in the abelianisation* and immune to an entire class of relations. Precisely, if $\kappa_\tau([w]) \neq 0$ then

$$[w] \notin [G, G],$$

so [w] has nonzero image in $G_{ab} = G/[G, G]$; and consequently w survives every coherence relation whose relator lies in the commutator subgroup [G, G].

Proof. $\kappa_\tau = \kappa \circ \tau_*$ is a homomorphism $G \to \mathbb{Z}_7$; since $\mathbb{Z}_7$ is abelian, $[G, G] \subseteq \ker \kappa_\tau$. If $\kappa_\tau([w]) \neq 0$ then $[w] \notin \ker \kappa_\tau$, hence $[w] \notin [G, G]$, so its image in G_{ab} is nonzero. For the immunity claim: imposing a relator $r \in [G, G]$ passes to a further quotient of G that leaves G_{ab} unchanged (r is already trivial in the abelianisation), so κ_τ — which factors through G_{ab} — is unaffected and continues to take the value $\kappa_\tau([w]) \neq 0$. Thus no commutator-valued relation can erase the charge.

This sharpens Detection Dominance from "the loop is nontrivial" to "the loop is *abelianisation-visible* and commutator-relation-proof." It complements the Residual Topology Theorem (§6.3) without being the same statement: §6.3 gives a *geometric* immunity (a global/charged loop is not a bookkeeping cycle, so the forced bookkeeping relations cannot kill it), while Corollary 10.7A gives an *algebraic* immunity (no relation whose relator lies in [G, G] can kill it, since the charge lives in the abelianisation). The two are independent guarantees pointing the same way; we do not claim the forced relators are themselves commutator-valued, only that both kinds of immunity shield a charged loop.

Lemma 10.7B (No Recovery) (irreversibility supplies confinement and permanence; essentialness inherited)

Suppose commitment discards a region D *and* D is **essential** in the post-commitment arena — that is, the boundary loop γ_D is nontrivial in $H_1(\mathcal{A})$, where $\mathcal{A}$ = (arena) \ D. Then γ_D is permanently non-bounding under admissible commitment-preserving motion.

Proof. A contraction of γ_D by admissible motion is a homotopy $H : S^1 \times [0,1] \to$ (arena), $H(\cdot,0) = \gamma_D$, $H(\cdot,1)$ constant — a filling disk whose image may a priori lie anywhere in the arena. Partition candidate fillings by whether the disk meets the discarded region D. *(i) The disk meets D.* Then the homotopy enters D; but irreversibility forbids admissible commitment-preserving motion from entering the discarded region, so no *admissible* filling can pass through D. Case (i) is excluded — not by typing H into $\mathcal{A}$ in advance, but as a conclusion: irreversibility removes exactly the through-D filling that would otherwise be available (and is available before commitment, when γ_D bounds the disk D itself). This is the native content. *(ii) The disk avoids D,* lying entirely in $\mathcal{A}$ = arena \ D. Such a disk exists iff γ_D bounds in $\mathcal{A}$, i.e. iff [γ_D] = 0 in $H_1(\mathcal{A})$; by essentialness [γ_D] ≠ 0, so no complement filling exists. Both cases are excluded, so γ_D admits no admissible filling — it is non-bounding, and permanently so, since irreversibility keeps D removed for all time.

Both premises are load-bearing, and which does what. The decomposition is exact. Irreversibility excludes the through-D filling — the genuinely native content — and essentialness excludes the complement filling — inherited topology. Neither alone suffices, and two counterexamples show it. *Drop essentialness:* on S^2, discarding a small disk D leaves ∂D bounding the complementary disk, so γ_D is contractible within $\mathcal{A} = S^2 \setminus D$ despite D being discarded; irreversibility forbids filling *through* D but says nothing about the far side. *Drop irreversibility:* the through-D filling is available and γ_D bounds in the full arena. So the lemma does **not** nativise non-bounding from bare discarding; it separates the two jobs and fixes exactly how much irreversibility contributes (confinement and permanence) and how much remains inherited (the essentialness — topological non-triviality — of D in the post-commitment complement).

Consequently "isolation," used below as the operative inherited premise, must be read in the **strong sense**: not merely that commitment discards D, but that D is *essential* in the post-commitment complement (γ_D ≠ 0 in $H_1(\mathcal{A})$). That essentialness is the inherited Gate-3 input, carrying whatever marker it holds in its home paper. What 10.7B genuinely buys is therefore narrower than a full nativisation: given strong isolation, the non-bounding and its permanence are native (irreversibility confines admissible motion to the complement and keeps it there); the topological weight of the claim sits visibly in essentialness, not hidden in the word "encloses."

Theorem 10.8 (Primitive-Fact Persistence) (conditional on refinement realisation)

Assume a primitive Fact isolates a discarded region D in the strong sense — D is discarded and *essential* in the post-commitment complement, γ_D ≠ 0 in $H_1(\mathcal{A})$ (the inherited input of Lemma 10.7B) — and that γ_D is refinement-realised by a closed word w (Def 10.4A, τ_* w = γ_D). Then the realised loop is non-bounding in $\mathcal{C}$_MS — G_ab ≠ 0, and RC_path does not close natively — for either of two reasons, both presupposing the strong isolation and realisation of the hypothesis and then drawing on different further inputs:

1. **Isolation protection (charge-independent).** γ_D is a *global* loop around the essential discarded region D, not a bookkeeping loop. By the Residual Topology Theorem (§6.3) the forced relations — cancellation, presentation, exchange — kill only bookkeeping cycles, so they cannot kill w. The only remaining killer would be a higher coherence (Conjecture 7.2); but by the No-Recovery Lemma (10.7B) γ_D is permanently non-bounding under admissible motion — through-D fillings excluded by irreversibility, complement fillings by essentialness. Hence w is not coherence-killable: [w] ≠ 0 in G_ab.

2. **Charge protection.** If additionally the closure phase is faithful and the $\mathbb{Z}_7$ closure charge is not globally inert, the minimal Fact cycle registers exactly one commitment's phase advance,

   $\kappa(\gamma\_D) = [\Delta\theta] \neq 0$ in $\mathbb{Z}_7$,

   so $\kappa\_\tau(w) \neq 0$ and, by Corollary 10.6, $[w] \neq 0$ in G_ab — and the surviving class carries nonzero κ-charge, identifying it with the closure class.

Proof. For (1): w is global by isolation, so by the Residual Topology Theorem it is not killed by the forced relations. For higher coherence, the transfer between complexes is explicit — τ is functorial, hence a chain map, so $\partial \circ \tau\_* = \tau\_* \circ \partial$; a 2-chain c in $\mathcal{C}$_MS with $\partial c = w$ transports to $\tau\_* c$ in Γ_vac with $\partial(\tau\_* c) = \tau\_* w = \gamma\_D$, so a filling of w *is* a filling of γ_D. Transported admissible chains are chains of the post-commitment arena (Γ_vac *is* that arena, with D removed), so the No-Recovery Lemma (10.7B) applies directly: γ_D is non-bounding there — through-D fillings excluded by irreversibility, complement fillings by essentialness — so no such c exists and w is non-bounding. For (2): Theorem 10.5 with the minimal-cycle phase relation supplies the nonvanishing of the charge, then Corollary 10.6. Either route gives $[w] \neq 0$ in G_ab = $H_1(\mathcal{C}\_MS)$, which by Theorem 6.1 is the failure of RC_path to close.

What the two protections do, separately. Charge protection (2) is the stronger conclusion — it makes the residue *the κ-class* — but it is hostage to the charge being globally active. Isolation protection (1) is the backstop precisely against that worry: even if one doubts κ is globally nonzero, a realised γ_D still cannot be erased, because its non-bounding is topological — D is essential in the complement (no far-side filling) and irreversible (no through-D filling) — rather than charge-theoretic. So the residue persists under (1) alone; under (1)+(2) it persists *and* is the Gate-3 class. This is the formal content of the general-reader claim that the trace cannot be smoothed away: undoing the loop would require either deforming through a region that commitment has irreversibly removed, or contracting an essential loop in the complement — and neither is available to admissible motion.

The hinge, unmoved. Persistence does not close the open question; it sharpens what the open question protects. Both protections are gated on refinement realisation (Def 10.4A): isolation makes γ_D unerasable *as a closure object*, and once realised as a word w, unerasable *in G_ab* — but it does not by itself establish that a reversible merge/split word transports onto γ_D. Outcome A (coherence, no native residue) remains available on exactly one condition: that γ_D is *not* refinement-realised. What Theorem 10.8 removes is the residual fear that, even granting realisation, ordinary coherence or a globally inert charge might still let the loop close. Neither can: a realised isolation loop is non-bounding for a reason no coherence relation and no inertness can touch.

The asymmetry, stated plainly. It is worth recording exactly which half of this is established and which conditional, because the two live on different sides of a seam. The chain *commitment discards D and D is essential → γ_D bounds neither through D (irreversibility) nor in the complement (essentialness) → γ_D is permanently non-bounding → any κ-charge on it persists* is a **closure-side result**, and Lemma 10.7B fixes precisely which parts are native: irreversibility nativises the confinement (no through-D filling) and the permanence, while the topological weight — that D is *essential*, so there is no far-side filling either — remains the inherited input. The remaining inherited premise is therefore not bare discarding but **strong isolation**: that commitment discards D *and* D is essential in the post-commitment complement ($\gamma\_D \neq 0$ in $H_1(\mathcal{A})$). Its standing is exactly the standing essentialness carries in its home papers; read here as *established (inherited) for the essentialness of D, native for the confinement and permanence that irreversibility then supplies*, not as independently proven end to end. With that caveat the trace's permanence is not in doubt on the closure side. What is conditional in a second, sharper sense is the transfer of that permanence to the reconstruction side: the RC residue lives in $G\_ab = H_1(\mathcal{C}\_MS)$, and for the closure-side permanence to become a permanent feature *there*, γ_D must be a class there — a reversible word w with $\tau\_* w = \gamma\_D$. Refinement realisation is the sole map-theoretic premise effecting that transfer, and it is exactly the premise the closure-side chain cannot supply, since it concerns an irreversible object while w is reversible. So the clean verdict is: **persistence forever is established on the closure side (essentialness inherited, confinement and permanence native), and established on the reconstruction side the moment refinement realisation is.** Without realisation, the permanence is real but is a property of the closure complex, not yet of merge–split holonomy — not because the RC side might fail to *notice* a trace that is there, but because, absent realisation, there is no RC-side object for permanence to be a property of.

10.4 The outcomes, under detection

- **Outcome A — coherence (Conjecture 7.2 holds).** Then by Corollary 10.6 the detection hypotheses must fail: γ_D is *not* refinement-realised (no reversible refinement loop transports onto it), or transport does not carry its charge. The native RC residue closes; the Gate-3 residue, though real, is then not seen by refinement motion, and the two programmes' holonomies are distinct.
- **Outcome B — detection (γ_D refinement-realised).** Then by Theorems 10.5–10.8, $G\_ab \neq 0$, Coherence Exhaustion fails, RC_path does not close natively, and the surviving native residue *persists* — unerasable by strong isolation (10.8(1)), and carrying nonzero κ-charge under faithful non-inert recording (10.8(2)). The RC residue and the Gate-3 residue are then established to *share charge*; whether they are the *same group* is the residual Φ-isomorphism question.

The two outcomes are mutually exclusive: Outcome A is precisely the failure of refinement realisation, Outcome B its success. And by Theorem 10.8 the question has now contracted to realisation *alone* — for once γ_D is realised, neither coherence nor a possibly-inert charge can erase it, so realisation is not merely necessary for the residue but, given isolation, sufficient for its persistence. The whole of the native RC question therefore turns on the single concrete property of Definition 10.4A: does the admissible transport map carry a reversible refinement loop onto the commitment cycle γ_D?

10.5 Why the recast matters

The Admissible Lift programme began from quantum reconstruction; the Gate-3 programme began from closure topology; for most of their development they appeared independent. The reductions converge them onto first homology, and the detection recast makes the cross-programme question concrete rather than comparative. It is no longer "are these two H^1 classes related?" but "does the reconstruction-side loop space register the closure-side charge that is already known to exist?" — answerable, in principle, by realising one cycle and computing one pairing.

---

11. Physical Interpretation and Programme Consequences

(interpretive throughout, conditional on residue survival)

The preceding sections are topological; their interest is physical. This section reads the residue as physics — but every statement in it is conditional on the residue actually surviving, and that survival rests on the one question the paper could not close: whether the primitive-Fact loop is refinement-realised (§10, §12). Read each "the residue means…" below as "if the residue survives, it would mean…". One asymmetry from §10.8 should be carried throughout: the *permanence* of the commitment trace is already a theorem on the closure side, unconditionally; what is conditional is only whether that permanence is also a feature of merge–split holonomy, which realisation decides. So the physics below stands on firmest ground read as a property of the closure substrate, and becomes a property of the reconstruction side exactly when realisation does.

With that frame fixed, the central question is one line:

Does irreversible commitment leave behind a permanent topological trace?

The analysis has shifted the burden of proof. The original position treated a surviving residue as a bare possibility. The primitive-Fact route makes a stronger conditional available: if primitive Facts exist and their cycles are refinement-realised, a nontrivial residue does not merely happen to survive — it survives *naturally*, protected by the irreversibility of commitment itself (§10.8). The consequences of that, should it hold, reach well beyond the RC programme.

11.1 Commitment as a source of topological memory

The most immediate consequence is ontological. A surviving residue would mean commitment is not purely local. The commitment event is local; the trace it leaves is not. The act changes the global topological structure of admissible motion, so that reality acquires — in a precise and limited sense — a *memory of commitment*.

That sense must be stated carefully, because it is not memory in the informational sense. No record of individual events is preserved, no chronology is stored, no database of history exists. What persists is a single global tally — the net closure charge, valued in $\mathbb{Z}_7$ — registering that commitment has occurred. And because the tally is valued mod 7, it is a *one-way* witness: a nonzero tally certifies that genuine commitment happened, but a zero tally does not certify that none did, since seven aligned commitments sum back to the identity (§8). The substrate does not remember events; it carries a running, cancelling count of net commitment, beneath which no finer record is kept.

This is nonetheless something new in the programme. Most of VERSF concerns closure, admissibility, commitment, and transport — the structure of how facts form. A surviving residue would add a different category: history itself, not as narrative but as persistent topology — the substrate's record not of *what* happened but *that* (net) commitment happened at all.

11.2 The emergence of phase

The Gate-3 programme already establishes the chain

ρ_rot → U(1) → continuum connection **(established in Gate-3)**,

by which the rotation-label cochain gives rise to a U(1) connection in the continuum description. A surviving residue makes this chain considerably more significant, because the surviving quantity is naturally *phase-like*: it is a holonomy, it accumulates around loops, and it is detected through transport — exactly the structures associated with quantum phase.

The interpretive step — and it is an interpretation, not a result — is then available: phase need not be a primitive feature of reality. It could instead be read as accumulated commitment residue, the continuum's name for the substrate's holonomy. In slogan form, conditional throughout,

phase ≈ memory of commitment.

Under this reading the role of phase in quantum theory becomes less mysterious: the substrate carries a record of commitment as holonomy, and the continuum description interprets that holonomy as phase. We mark the identification conjectural; what the paper supports is the weaker, structural claim that the surviving quantity has the formal character of a phase.

11.3 Implications for the Born rule

The Born-rule programme repeatedly meets the same conceptual question: why should amplitudes carry phase, and why should interference survive? A persistent κ-class offers a candidate answer, again conditional. Facts leave a topological trace; the trace is a holonomy; holonomy presents as phase; and interference is the route-dependence of that accumulated phase — the same route-dependence this paper calls RC_path.

The consistency demanded by §11.1 sharpens the last step. Because no per-event history is stored, interference cannot be read as interaction between stored individual histories. It is better read as the path-dependence of a single accumulated charge: two routes between the same endpoints differing by a closed-loop holonomy, which is precisely a nonzero element of G_ab. On this reading interference and the RC residue are the same phenomenon seen twice. The residue would thereby supply a physical interpretation for structures that appear only formally in the ODG and OIP programmes — the companion strands in which amplitude and phase structure enter as formal apparatus rather than as derived physics.

11.4 Global information

The residue is not stored at a point, not attached to a particle, not associated with any local observable; it is measurable only around loops. That places it among the global topological quantities — the kind familiar from topological quantum theories, where information is encoded globally rather than locally. A surviving residue would therefore suggest a route toward understanding why quantum information so often appears non-local in character.

The claim is deliberately modest. It is not that non-locality has been explained; it is that the substrate may already contain the kind of global structure from which such behaviour could naturally emerge — a place for the phenomenon to come from, not a derivation of it.

11.5 Relation to gravity

The implications for gravity are weaker, and worth stating as such so the section does not overreach. The gravity programme proceeds through commitment → fact density → fact momentum → geometry; the residue does not alter that chain. Instead it introduces a second, parallel branch from the same root: commitment → persistent topological memory. The residue is thus more closely related to information and phase than to curvature, and its strongest consequences are likely to lie on the quantum side of the programme rather than the gravitational side. The two branches share only their origin in commitment.

11.6 The first history variable

The deepest consequence may be conceptual. If the residue survives, the substrate possesses a quantity whose value depends on the existence of irreversible commitments — neither geometric nor dynamical in the usual sense, but *historical*. The universe would retain a record that net commitment occurred: not where, not when, not in what order, and not even that any *particular* commitment occurred — only the net, cancelling total valued in $\mathbb{Z}_7$.

Stated at that strength, this would be the first genuine *history variable* identified within VERSF — the programme's first quantity that exists because the past is fixed, rather than because of any configuration in the present. And it rests squarely on the open hinge: the variable is a closure-side reality already, since the permanence is established there (essentialness inherited, confinement and permanence native, §10.8), and it becomes a variable of the reconstruction substrate the moment refinement realisation is established, and not before. If that step holds, the small unerasable tally of the general-reader summary is not a metaphor but the substrate's first memory — the record, kept in sevens, that reality has committed.

11.7 Why the residue's location was inevitable

It is worth ending on why the residue, should it exist, was bound to live exactly where this paper finds it — because an inevitability argument is more trustworthy than a surprise, and because what is inevitable here is the *location*, which is proven, not the *existence*, which is not. The chain is short and each link is established earlier in this paper or its companion:

1. **Commitment isolates, essentially.** An irreversible Fact discards a region D, and D is *essential* in the post-commitment complement — γ_D nontrivial there (strong isolation, §10.7B). This is the inherited input; its essentialness, not the bare discarding, is the load-bearing part.
2. **Essential isolation creates protected cycles.** Given essentialness, γ_D bounds nothing in the complement; given irreversibility, no admissible motion can fill it through the discarded D either (No-Recovery Lemma, 10.7B — irreversibility native, essentialness inherited). So commitment manufactures non-bounding loops as a matter of course — the topological weight supplied by essentialness, the permanence by irreversibility.
3. **Protected cycles are homology.** A non-bounding loop is by definition a nontrivial class in first homology; route-dependence of any charge it carries is an H^1 phenomenon (§2.2), and nothing of lower or — by Proposition 2.3 — higher degree can carry that charge.
4. **First homology is G_ab.** Reversibility makes the merge–split structure a groupoid, connectedness collapses it to one group, and H_1 is its abelianisation (Theorem 3.5). So the only home a protected closure charge can have is G_ab.

Read forward, the chain says: *given that commitment isolates, a persistent residue is not an exotic possibility to be hunted for but the expected shadow of irreversibility, and it can live in one place only.* The residue is thus demoted from a surprise to a structural near-corollary of commitment — with two honesties kept in view. The existence of a nonzero residue still waits on refinement realisation (the chain establishes where, not whether), and step 1 still imports isolation. Inevitability here is the inevitability of the address, not of the tenant: if anything persists, it was always going to be a class of G_ab, detected as a κ-charge, and nothing else.

---

12. What Remains Unknown

The paper proves the localisation and reduces the residue to one group; it does not compute that group, and it does not settle the cross-programme identification. The open questions are now sharply posed.

- **Coherence Exhaustion (Conjecture 7.2) — the central open question.** Whether the forced relations plus associativity/coassociativity *exhaust* the admissible coherences. If they do, $G\_ab = 0$ and RC_path closes; if not, a residual class survives. This is a definite presentation question for G, not a conceptual doubt. **(conjectural)**

- **Existence of a finite-order / sevenfold cycle (§8).** Whether $\mathcal{C}$_MS contains a nontrivial torsion holonomy, and specifically an order-7 one. This fixes whether Outcome B is κ-compatible. **(open)**
- **The H^{2+} degree question (§2.2), now narrowed.** Whether $\mathcal{C}$_MS carries seal-bearing higher cohomology at all. By Proposition 2.3 any such class is *charge-free* — H^{2+} cannot contribute to the closure-charge lift, to κ-detection, or to RC_path as defined — so this no longer qualifies the residue the paper is about. It survives only as the bare-mathematical question of whether some other, non-closure-charge seal lives in higher degree. Inherited from the companion's §9.8; it presupposes topology and asks only about degree. **(open, charge-free)**
- **Refinement realisation of γ_D (§10, Def 10.4A) — the now-sole decisive question.** Whether the admissible transport map carries some closed reversible refinement word onto the commitment cycle γ_D. By Theorems 10.5–10.8, *if* it does, then $G\_ab \neq 0$, Coherence Exhaustion fails, and the native residue persists — unerasable by strong isolation (D essential, §10.7B), and carrying nonzero κ-charge under faithful non-inert recording. Persistence (10.8) shows realisation is, given strong isolation, not just necessary but sufficient: neither coherence nor an inert charge can then close the loop. The crux is the commitment/pre-commitment category gap — γ_D irreversible and downstream, the realising word reversible and upstream. **(open, decisive)**
- **The Γ_vac identification / full Φ (§10).** Whether the comparison map $\Phi : H^1(\mathcal{C}\_MS; A) \to H^1(\Gamma\_vac; \mathbb{Z}_7)$ is an *isomorphism* on the relevant classes — equivalently whether G_ab and $H_1(\Gamma\_vac)$ are the same holonomy group, not merely sharing the detected κ-charge. Detection (above) gives a nonzero homomorphism; this asks for the full identification. **(open)**
- **The non-abelian remainder (§1).** The closure-charge lift is settled entirely by G_ab (its monodromy is $\mathbb{Z}_7$-valued, hence factors through the abelianisation). A hypothetical *non-abelian-valued* lift would instead require $G = 1$, not merely $G\_ab = 0$; whether the commutator subgroup [G, G] is trivial is a separate question, distinct from all of the above — and in particular distinct from the H^{2+} degree question, which concerns cohomological degree rather than [G, G]. It carries no closure charge and is invisible to κ, so it bears on RC_path only under the non-abelian reading the paper does not adopt. **(open, charge-free)**

Two prior structural facts remain presupposed from the companion and worth confirming independently: that $\mathcal{Q}$_full is connected as established (used throughout to give $\mathcal{C}\_MS \simeq BG$), and that the elimination programme of §2.1 is complete as inherited. A third, load-bearing for persistence specifically, is the **essentialness of the discarded region D** (γ_D nontrivial in the post-commitment complement, §10.7B): this is the topological weight of strong isolation, inherited from the closure papers, and the persistence conclusion is only as secure as it is.

The remaining work is mechanical once the data are specified: the admissible merge generators, the admissible split generators, the complete coherence-relation set, the exchange relations, and the presentation relations determine the chain complex, and one then constructs

$\partial_1 : C_1 \to C_0$, $\partial_2 : C_2 \to C_1$, $H_1(\mathcal{C}\_MS) = \ker(\partial_1) / \mathrm{im}(\partial_2) = G\_ab$.

If H_1 vanishes, the native residue closes; if it survives, the surviving class is the unique native RC residue, and its order is compared with κ.

---

13. Conclusion — The Endgame

The reduction programme can now be stated as a single ladder, with the status of each rung explicit.

1. **Eliminate local candidates** — metric, label, ordering, history, local-combinatorial. **Completed (established).**
2. **Eliminate component-seals** — $H^0(\mathcal{Q}\_full) = \mathbb{Z}$, the refinement quotient is connected; reachability (RC_reach) holds natively. **Completed (established).**
3. **Localise the remaining residue** — every surviving native residue is a route-dependence (holonomy), hence a class of $H^1(\mathcal{C}\_MS)$. **Completed (proven, this paper).**
4. **Sharpen the object (Residue Reduction Theorem, §3.5)** — reversibility makes MS(M) a groupoid; connectedness gives $\mathcal{C}\_MS \simeq BG$, so the native RC residue *is* $G\_ab = H_1(\mathcal{C}\_MS)$, with $H^1(\mathcal{C}\_MS; A) = \mathrm{Hom}(G\_ab, A)$; it survives iff G is not perfect. **Completed (proven, this paper).**
5. **Derive the forced relations** — cancellation (FP1), presentation coherence (IA), exchange (independence) are forced, and any survivor is genuinely topological (Residual Topology Theorem). **Completed (proven).**
6. **Decide higher-coherence exhaustion** — whether the forced relations collapse G_ab to zero. **Open (conjectural)** — and, by rung 8, falsified if detection holds.
7. **Compute $H_1(\mathcal{C}\_MS) = G\_ab$** from the coherence presentation. **Open** — but no longer the only route to the residue branch, given rung 8.
8. **Refinement-realise γ_D (Theorems 10.5–10.8).** Exhibit a reversible refinement word transporting onto the commitment cycle γ_D. If it exists, the loop persists — unerasable by strong isolation, κ-charged under faithful recording — so $G\_ab \neq 0$ and the native residue is real, deciding rungs 6–7 without the full coherence computation. **Open, decisive.**

The remarkable feature of this progression is that the original RC question has *disappeared*. It is no longer "does some hidden obstruction exist somewhere in refinement space?" nor even "why should reality be connected?" The merge–split programme began as an attempt to determine whether admissible refinement motion possesses a surviving route-dependence residue at all. The Gate-3 programme has since changed the terms: it independently identified a closure-transport residue $\kappa \in H^1(\Gamma\_vac; \mathbb{Z}_7)$, carried by protected primitive-Fact cycles. What remains is therefore not whether κ exists — it does, on the closure side — but whether admissible refinement motion *detects* it.

The honest terminus on RC alone is unchanged and still stands: the negative half is proven down through H^0 and localised in H^1; the positive default (RC_path closes) rests on Coherence Exhaustion, which would make G perfect — most simply trivial, with $\mathcal{C}$_MS contractible. What the detection recast adds is a second, sharper route to the *opposite* verdict. By Theorem 10.5 and Corollary 10.6, if the primitive-Fact loop γ_D is a genuine refinement loop carrying its Gate-3 charge, then G_ab ≠ 0, Coherence Exhaustion is false, RC_path does not close natively, and the surviving native class carries nonzero κ-charge. The two routes are the two faces of one structural fact: coherence (G perfect, residue closes, κ undetected) versus detection (γ_D charged, residue survives, κ detected) are mutually exclusive and jointly cover the question.

So the analysis no longer suggests merely that the reconstruction residue *resembles* the Gate-3 residue. Under detection it would *be* the Gate-3 residue, seen from the reconstruction side — sharing charge by Corollary 10.6, with only the full Φ-isomorphism (whether the two holonomy groups coincide, not merely share a charge) left to settle whether they are literally one group. The remaining question is therefore not whether κ exists, but whether the merge–split holonomy group and the closure-transport holonomy group are different presentations of the same underlying structure.

The programme has bifurcated. The single open question — is G_ab trivial? — is now approachable by two routes that resolve it in opposite directions and cannot both succeed:

- **Route A — prove Coherence Exhaustion (§7).** Establish that the forced and higher relations exhaust all admissible composition cycles. Then G_ab = 0 (G perfect): the H^1/detectable residue closes and κ is undetected by refinement motion. This does *not* give contractibility — a perfect nontrivial G has $H_1(BG) = 0$ but generally nonzero Schur multiplier $H_2(BG)$, which is exactly the H^{2+} mechanism left open in §2.2. Contractibility, and with it the closing of every degree, belongs only to the sub-case G = 1.
- **Route B — prove refinement realisation of γ_D (§10).** Exhibit a closed reversible refinement word that transports onto the commitment cycle γ_D. Then by Theorems 10.5–10.8, G_ab ≠ 0, RC_path does not close natively, and the surviving native class *persists* — unerasable because, under strong isolation, the loop bounds neither through the irreversibly-removed region nor in its complement — and carries the κ-charge under faithful non-inert recording, independently of any coherence computation.

The two routes are mutually exclusive: A proves G_ab = 0, B proves G_ab ≠ 0, and G_ab has one definite value, so at most one route can succeed, and which one is fixed by the truth about G. They are exhaustive of the *verdict* — G_ab either vanishes or it does not — with one honest asymmetry worth recording: Route A, if it succeeds, closes the H^1/detectable residue completely (no κ-class from any loop), though it leaves the H^{2+} degree question (§2.2) open unless G is moreover trivial; Route B exhibits one sufficient witness for the residue, not the only conceivable one (G_ab could be nonzero through some loop other than γ_D). So success of A closes the detectable residue; success of B establishes the residue and its κ-charge while leaving the full holonomy-group identification (Φ) as the one further step. Either completion ends the central RC_path question; neither requires the other.

That is the strongest reduction presently available. The final uncertainty of the programme has been compressed to a single structural question about a single named cycle: is the primitive-Fact loop a charged refinement loop? If it is, the reconstruction-side residue and the closure-side residue share one charge, and the last open question is whether they share one group.