The Refinement Realisation Criterion

Transport-Surjectivity, Primitive-Fact Cycles, and the Final Gate-3 Hinge

Keith Taylor — VERSF Theoretical Physics Programme

General Reader Summary

What this paper settles, and what it leaves open

The whole question of whether reality keeps a permanent trace of irreversible events now comes down to one precise, almost mechanical check — and this paper turns that vague-sounding question into an exact one, proves the part that can be proved today, and marks the part that cannot.

Earlier work in the programme reached a striking position: *if* a certain loop (the trace a committed, irreversible event leaves behind, written γ_D) can be reached by ordinary reversible refinement motion, then reality keeps a permanent topological memory of commitment — and the quantum phase that physics usually assumes would acquire a physical origin. Everything hung on one word: *reachable*. This paper is about pinning down exactly what "reachable" means and what it would take to decide it.

The first thing the paper does is separate two questions that are easy to confuse. One is whether the trace γ_D *survives* — whether, once an irreversible event carves it out, it can never be smoothed away. That was settled earlier: it survives. The other is whether ordinary reversible motion can *reach* it. These are different questions about different things, and the paper proves they are different: survival does not give reachability for free. The whole open problem lives in that gap.

The paper then makes "reachable" precise. Reaching γ_D means a certain map — admissible transport — carries some reversible loop onto it. Whether it does turns out to split cleanly into two stages. The first stage is a *necessary* check that can be done by ordinary (if laborious) bookkeeping: write down two lists of basic loops, record where the transport map sends each, assemble the numbers into a grid, and test whether γ_D's numbers can be built from that grid using whole numbers only. If they cannot, the answer is *no* — reality keeps no such trace — and the question is closed, finitely and decisively. The paper specifies this computation exactly, down to the standard algorithm that runs it, while being honest that four pieces of input data are not yet in hand to run it numerically.

But passing that first check is *not* enough to say *yes*. The paper is careful about this, because it is the place every loose version of the argument goes wrong. The bookkeeping check works at the level of homology — a coarse accounting that asks whether γ_D's class *could* be built up from transported pieces. Actually reaching it requires a *single reversible loop* to land on it exactly, and that is a strictly harder thing, made harder by the very reversibility that distinguishes refinement motion from the irreversible event it is trying to reach. So the bookkeeping check can *refute* the whole picture but cannot *confirm* it; confirmation needs a further result the paper names but does not prove.

Finally, the paper records a hint about which way the answer may fall. Several independent strands of the programme each turn up a *single* persistent direction — one cycle, one mode, one class — rather than a sprawl of them. If these are all the same direction (which is conjectured, not proved), then the space the transport map can possibly reach is just one line, and the necessary check collapses to asking whether γ_D points along it. The strongest version of this hint is that the refinement framework's own topology is rank one — a single loop — which, if it carries over to the transport setting, would by itself force the reachable space to be at most one line. The paper does not leave that "if it carries over" as a hope: it sets out exactly the test the carry-over must pass — that every reversible loop be built from the one basic loop (and, a subtler point it is careful about, that this basic loop genuinely has infinite extent rather than secretly closing up after a few steps). Two honesties are kept even here: "at most one line" includes the possibility of *no* line at all — a reachable space of zero, which would be the outright negative answer, reality keeping no trace — and the hint does not decide between the single line and nothing. The paper keeps the whole thing firmly at the level of suggestion: it is where the evidence points, not what the paper proves, and the evidence points at "one line or none" before it points at "one line."

In short: the question of reality's memory of commitment is now a definite mathematical target with a known shape — a computation that can say *no* outright, and a named further theorem that would be needed to say *yes* — plus a glimpse, no stronger than a glimpse, that the target may be as simple as a single line.

Abstract

Recent Gate-3 and Reversible-Connectedness results reduce the question of a native closure residue to a single unresolved issue: whether the primitive-Fact cycle γ_D is *refinement-realised* — whether a reversible merge–split loop transports onto it.

This paper formalises that question and locates its difficulty precisely. Refinement realisation is, by definition, the membership of γ_D in the image of admissible transport restricted to reversible loops; the content of the paper is not that restatement but what surrounds it. We prove a **homological obstruction** — a necessary condition, $[\gamma_D] \in \tau_*(H_1)$, that is checkable by a homology computation without solving the full realisation problem — and we isolate, as the paper's central honesty, the **gap between the homological image and the loop-level image**:

membership in homology is necessary but *not sufficient* for realisation, because a class hit by a chain need not be hit by a single reversible loop. That gap, sharpened by the reversibility restriction on the transport map, is the genuine open problem the programme now faces.

We then state what transfers if realisation holds — not "every closure-side invariant," but every invariant *natural under transport*, with the κ-charge and closure holonomy as the instances that qualify. The result is to compress the native-residue programme to one concrete, correctly-scoped target:

Is [γ_D] in the loop-level image of reversible admissible transport?

A positive answer transfers κ-charge, closure holonomy, and the Reversible-Connectedness residue into admissible refinement motion, making phase-as-memory a viable physical interpretation. A negative answer — most cleanly detected when $[\gamma_D] \notin \tau_*(H_1)$ — confines the residue permanently to closure topology. The paper does not decide the question; it establishes the criterion, the one checkable obstruction, and the precise location of what remains.

Epistemic markers: ***(established)*** */* ***(inherited)*** *for results from prior VERSF papers;* ***(proven)*** *for results proved here;* ***(definition)*** *for definitional content;* ***(open)*** *for what remains undecided.*

Table of Contents

---

1. Introduction

The Gate-3 programme and the Reversible-Connectedness (RC) programme began from apparently distinct questions. Gate-3 investigated whether irreversible commitment leaves behind protected topological structure. The RC programme investigated whether admissible refinement motion carries any native route-dependence once all eliminable structure has been removed. Initially these appeared unrelated; recent results have converged them.

The persistence branch of Gate-3 identifies a protected primitive-Fact cycle γ_D carrying nontrivial closure holonomy — permanently non-bounding on the closure side, given irreversibility and the essentialness of the discarded region. The RC reduction identifies the surviving native residue, if one exists, with transport structure visible to admissible refinement motion, and isolates a single decisive hinge. Together they imply that the entire native-residue question reduces to one issue:

Does admissible reversible transport realise the primitive-Fact cycle γ_D?

The purpose of this paper is to formulate that question precisely and determine exactly what must be proved to answer it. **The paper does not establish refinement realisation.** It does three things instead: it states the realisation condition as a definite image-membership (§3); it proves the one checkable necessary condition — a homological obstruction (§4); and it identifies, as the genuine open problem, the gap between that homological condition and the loop-level condition realisation actually requires (§6). The honest summary is that this paper converts a vague question into a sharp one and proves the part that is provable now, while naming precisely the part that is not.

A word on what this paper is and is not. It is a *criterion* paper: its value is in localisation and in separating the checkable from the open, not in new closure topology or new transport constructions. Where earlier drafts of this material risked labelling a definitional restatement as a theorem, this version is careful to mark definitional content as definition, the genuine necessary condition as the theorem, and the residual difficulty as open. That discipline is the point — a criterion paper earns its place only by being exact about how much it has and has not done.

---

2. Persistence and Realisation Are Distinct

A distinction must be maintained throughout, because conflating its two halves is the easiest way to overstate the result.

Closure persistence is a property of γ_D *within the closure complex*: that γ_D is permanently non-bounding there, given irreversibility and essentialness (the Gate-3 persistence result, inherited). It asks: *does γ_D survive?*

Refinement realisation is a property of the *transport image of reversible loops*: that some reversible merge–split loop transports onto γ_D. It asks: *can reversible transport reach γ_D?*

These are properties of different objects — one of a cycle inside C_cl, the other of the image of a map out of the reversible loop space — and the programme's recent reductions show that the native-residue question turns on the second, not the first. Persistence is established (modulo its inherited inputs); realisation is open.

Lemma 2.1 (Persistence does not entail realisation) (proven)

Closure persistence of γ_D does not imply refinement realisation of γ_D. Formally: γ_D being non-bounding in C_cl does not entail $[\gamma\_D] \in \mathrm{Im}(\tau\_loops)$.

Proof. Persistence asserts $[\gamma\_D] \neq 0$ in $H_1(C\_cl)$ — a statement about C_cl alone, quantifying over fillings within the closure complex. Realisation asserts $[\gamma\_D] \in \mathrm{Im}(\tau\_loops)$ — a statement about the map τ and its reversible-loop domain. The first places γ_D in the non-trivial part of $H_1(C\_cl)$; the second places it in a particular subset of $H_1(C\_cl)$ (the loop-level transport image) that is determined by τ and Γ_MS, not by C_cl. Non-triviality in $H_1(C\_cl)$ carries no information about membership in that image: a class can be non-zero in $H_1(C\_cl)$ while lying outside $\mathrm{Im}(\tau\_loops)$, since nothing about being unfillable-in-C_cl constrains whether a reversible loop of Γ_MS transports onto it. Hence persistence does not entail realisation. ■

This is the load-bearing direction, and the paper's entire spine rests on it: persistence is secured, realisation is open, and the gap between them is the open question precisely because this entailment fails. We state it as a proven lemma rather than leaving it at observation strength, because the paper relies on it everywhere and the argument — non-bounding is about C_cl, image-membership is about τ's domain — is complete.

Observation 2.2 (Full independence is not claimed) (open)

The *converse* direction — that realisation does not entail persistence — and hence full logical independence, we do **not** prove. A genuine independence claim needs witnesses in both directions, and while Lemma 2.1 secures one direction by a general argument, the other would require exhibiting a realised-but-non-persistent cycle, which we do not. We therefore claim exactly what is proved: persistence $\not\Rightarrow$ realisation (Lemma 2.1, the direction the programme uses); realisation $\not\Rightarrow$ persistence and full independence remain open for want of the harder witness. (A schematic witness in the persistent-but-unrealised direction — a closure complex with non-bounding γ_D that no reversible loop maps onto — would additionally make Lemma 2.1 *sharp* by showing the non-entailment is realised, not merely possible; constructing one is left open.)

---

3. The Transport Map and the Realisation Condition

Let Γ_MS denote the merge–split transport complex and C_cl the closure complex. Let

τ_loops : Loop_rev(Γ_MS) $\longrightarrow$ Z_1(C_cl)

denote admissible transport from *reversible* merge–split loops to closure 1-cycles. The restriction of the domain to **reversible** loops, written Loop_rev, is not cosmetic: it is the source of every difficulty below, because the image of τ on reversible loops can be strictly smaller than its image on all chains. The object of interest is the primitive-Fact cycle γ_D.

A choice must be fixed at the outset, because it changes what is being asked. Realisation can be demanded **on the nose** (cycle equality, τ_loops(w) = γ_D in Z_1(C_cl)) or **up to homology** (class equality, [τ_loops(w)] = [γ_D] in H_1(C_cl)). The RC framing is class-level, and we adopt it as primary: realisation up to homology is what residue transfer requires, since the transported invariants (κ-charge, holonomy) are themselves homological. We note where the on-the-nose version differs, as it is strictly stronger and would make the obstruction below necessary for a harder target.

Definition 3.1 (Refinement realisation) (definition)

The primitive-Fact cycle γ_D is **refinement-realised** if there exists a reversible merge–split loop w with

[τ_loops(w)] = [γ_D] in H_1(C_cl)

(class-level; the on-the-nose variant replaces this by τ_loops(w) = γ_D in Z_1(C_cl)).

Equivalently — and this is a restatement of the definition, not a theorem — γ_D is refinement-realised iff [γ_D] lies in the loop-level image

[γ_D] $\in$ Im(τ_loops) (as classes),

since "[γ_D] $\in$ Im(τ_loops)" unfolds, by the definition of image, to "$\exists$ reversible loop w with [τ_loops(w)] = [γ_D]." We flag this explicitly because it is tempting to dress the equivalence as a criterion theorem; it is not — it is the definition of image read twice. The content of the paper begins at the next section, where a *checkable* condition that does not presuppose knowing the image is established.

---

4. The Homological Obstruction — the Paper's Result

The transport map induces a homomorphism on first homology,

τ_* : H_1(Γ_MS) $\longrightarrow$ H_1(C_cl), whose image $\tau_*(H_1(\Gamma\_MS))$ is a *computable* object — a subgroup of $H_1(C\_cl)$ determined by homology, not by the loop space. This is the lever the paper actually provides.

Theorem 4.1 (Homological obstruction) (proven)

If γ_D is refinement-realised, then

$[\gamma\_D] \in \tau_*(H_1(\Gamma\_MS))$.

Proof. Realisation gives a reversible loop w with $[\tau\_loops(w)] = [\gamma\_D]$ in $H_1(C\_cl)$. A loop w determines a class $[w] \in H_1(\Gamma\_MS)$, and transport commutes with the passage to homology, so $\tau_*([w]) = [\tau\_loops(w)] = [\gamma\_D]$. Hence $[\gamma\_D] \in \tau_*(H_1(\Gamma\_MS))$. ■

Corollary 4.2 (Homological disqualifier — the operational payoff) (proven)

If

$[\gamma\_D] \notin \tau_*(H_1(\Gamma\_MS))$,

then γ_D is **not** refinement-realised.

Proof. Contrapositive of Theorem 4.1. ■

This corollary, not the definition of §3, is the paper's operational result. It gives a way to *kill* refinement realisation — and with it the native residue, the κ-transfer, and the phase-as-memory interpretation — by a single homology computation, without ever characterising the loop-level image. One computes $\tau_*(H_1(\Gamma\_MS)) \subseteq H_1(C\_cl)$, locates $[\gamma\_D]$, and checks membership. If $[\gamma\_D]$ falls outside the subgroup, the entire downstream programme is settled negatively, finitely, and without any loop-space analysis. That is a genuine reduction: the hardest-to-attack object (the image of τ on reversible loops) is replaced, *for the purpose of a negative verdict*, by a homology-group computation.

What the corollary cannot do is deliver a *positive* verdict, and the next section makes the necessary computation explicit before §6 explains why even passing it is not enough.

---

5. Computing the Homological Image

The homological obstruction is not merely conceptual. Once bases for $H_1(\Gamma\_MS)$ and $H_1(C\_cl)$ are chosen, it becomes an integer linear-algebra problem: compute the transport matrix M_τ and test whether the coordinate vector of $[\gamma\_D]$ lies in its column span over $\mathbb{Z}$. This section sets out that procedure, so that the moment the four missing inputs are supplied (listed at the end), the check can be run mechanically.

5.1 The procedure

Choose generators. Pick cycle generators for the two first-homology groups:

$a_1, \dots, a_r$ for $H_1(\Gamma_{MS})$, $b_1, \dots, b_s$ for $H_1(C_{cl})$.

Compute the transport images. Express each transported generator in the target basis — each $\tau_*(a_i)$ is an integer combination of the b_j:

$\tau_*(a_i) = m_{1i}\, b_1 + m_{2i}\, b_2 + \cdots + m_{si}\, b_s$.

Assemble the transport matrix. Collect the coefficients into an $s \times r$ integer matrix

$M_\tau = (m_{ji})$,

whose i-th column is the coordinate vector of $\tau_*(a_i)$. Then the homological image is exactly the column span (over $\mathbb{Z}$) of this matrix:

$\tau_*(H_1(\Gamma_{MS})) = \mathrm{Col}_{\mathbb{Z}}(M_\tau)$.

Run the test. Let g be the coordinate vector of $[\gamma_D]$ in the basis $b_1, \dots, b_s$. The homological obstruction (§4) is then the solubility of

$M_\tau x = g$

in integers $x \in \mathbb{Z}^r$:

- **No integer solution** $\Longrightarrow$ $[\gamma_D] \notin \tau_*(H_1(\Gamma_{MS}))$ $\Longrightarrow$ by Corollary 4.2, refinement realisation **fails**, and the programme is settled negatively by this computation alone.
- **An integer solution exists** $\Longrightarrow$ $[\gamma_D]$ passes the homological obstruction — it is in the homological image — but the loop-level realisability problem remains open (§6). Passing the test is necessary, not sufficient.

5.2 Why the solve must be over $\mathbb{Z}$, not $\mathbb{R}$

The solubility of $M_\tau x = g$ must be tested over the integers, and this is not a technicality in a programme where torsion is load-bearing. A real or rational solution x does not correspond to a homology class — homology classes are integer combinations of generators, so a vector with non-integer entries names nothing in H_1. More sharply, the gap between real and integer solubility is exactly where torsion lives: a class can be $\mathbb{R}$-reachable yet $\mathbb{Z}$-unreachable precisely because of finite-order phenomena, and the $\mathbb{Z}_7$ structure pervading this programme is finite-order throughout. The correct procedure is therefore the integer one: reduce M_τ to **Smith normal**

form, $U M_\tau V = D$ with U, V invertible over $\mathbb{Z}$ and $D = \mathrm{diag}(d_1, \ldots, d_k, 0, \ldots)$ the invariant factors, and read off integer solubility of $M_\tau x = g$ from D and the transformed target U g in the standard way — solvable iff each d_t divides the corresponding entry of U g and the entries beyond the rank vanish. A real-linear-algebra solve (Gaussian elimination over $\mathbb{Q}$) would give the wrong verdict exactly on the torsion cases the programme most cares about.

5.3 The four inputs the computation needs

The procedure is fully specified; what is not yet in hand is its data. To run it numerically requires:

1. **generators of $H_1(\Gamma\_MS)$** — the a_i;
2. **generators of $H_1(C\_cl)$** — the b_j;
3. **the action of τ on each generator** — the integer coefficients m_{ji}, i.e. the matrix M_τ;
4. **the coordinate vector of $[\gamma\_D]$** — the target g.

We do not supply these here, and we do not pretend to: without them no honest numerical verdict is possible. What this section establishes is that the obstruction is *computationally definite* — once (1)–(4) are provided, the stage-one verdict is a finite, mechanical Smith-normal-form computation with no remaining conceptual freedom. The paper thereby converts "check whether $[\gamma\_D]$ lies in the transport image" from an aspiration into a procedure waiting on four named inputs, and isolates exactly which inputs the closure and merge–split analyses must still supply.

A caution carried over from §4 and sharpened: this procedure decides only the *necessary* condition. An integer solution to $M_\tau x = g$ certifies $[\gamma\_D] \in \tau_*(H_1)$, the homological image — not $[\gamma\_D] \in \mathrm{Im}(\tau\_loops)$, the loop-level image. The next section is about that remaining gap, which no enlargement of this linear-algebra computation can close, because it is not a homological question.

5.4 The image as a horizon

It is worth recording what kind of object $\tau_*(H_1(\Gamma\_MS))$ is, because it is more than a test one runs against γ_D.

Proposition 5.4 (Maximal homological residue) (proven)

The subgroup $\tau_*(H_1(\Gamma\_MS)) \subseteq H_1(C\_cl)$ is the **largest homological residue admissible transport can detect**: every closure class transport can register at the level of homology lies in it, and any class outside it is invisible to admissible transport, regardless of loop-level considerations.

Proof. The statement has two halves, and they are not of equal depth. **Homological reach (tautological).** Anything transport can present to refinement motion is, by construction, the image under τ of something in its domain; passing to homology, every transport-visible class is τ_* of a class in $H_1(\Gamma\_MS)$, hence lies in $\tau_*(H_1(\Gamma\_MS))$. This much is the definition of image, and establishes only the *homological* statement: a class outside $\tau_*(H_1)$ is not τ_* of anything, so nothing transports onto it *at the level of homology*. **Upgrade to all loops (via Theorem 4.1).** The strong claim — that such a class is unreachable by *any loop*, not merely homologically — does not follow from the tautology; it requires the one-directional implication of Theorem 4.1, that loop-realisation implies homological membership (a realising loop induces a homology class, and τ_* of that class is the loop's transported class). Contrapositively: if $[c] \notin \tau_*(H_1(\Gamma\_MS))$, then no reversible loop can realise it, because realisation would force [c] into the homological image. So the homological exclusion transfers *downward* to a loop-level exclusion — and this is the entire content beyond the tautology. ■

The horizon, and why it is loop-decisive in only one direction. This is the place §5 and §6 must be reconciled rather than left in apparent conflict, because §6's whole thesis is that homology *underdetermines* loop-level — yet Proposition 5.4 says a homological verdict is loop-level-final. Both are true, and the reconciliation is the asymmetry that is this paper's subject: **an upper bound transfers downward but not upward.** Homological reachability is a *necessary* condition for loop reachability (Theorem 4.1, loop ⟹ homology), and never the reverse. Therefore:

- **Outside the horizon is final.** If $[\gamma\_D] \notin \tau_*(H_1)$, it is unreachable by any loop, ever — the negative verdict transfers from homology down to loops, because nothing loop-reachable can sit outside the homological image. This is the one direction in which a homological computation is loop-level-decisive.
- **Inside the horizon is not.** If $[\gamma\_D] \in \tau_*(H_1)$, nothing about loop-realisation follows — the positive direction does *not* transfer upward from homology to loops, and that non-transfer is exactly the §6 gap.

So $\tau_*(H_1)$ is a genuine *horizon* in the precise sense that it bounds transport's reach from outside: it is an upper bound, and an upper bound kills the positive case nowhere (a class inside may still be loop-unreachable) while securing the negative case everywhere (a class outside is loop-unreachable, full stop). Proposition 5.4 is therefore not in tension with §6; it is the sharp statement of §6's asymmetry — the homological image decides reachability negatively and only negatively.

Read against γ_D, the §4 obstruction is exactly this one-directional verdict: asking whether $[\gamma\_D] \in \tau_*(H_1)$ is asking whether γ_D is within transport's homological field of view at all. If it is not, γ_D is permanently beyond any loop — not because Proposition 5.4's tautology says so, but because Theorem 4.1 forbids a loop from reaching outside the homological image. That is a sharper *negative* verdict than Corollary 4.2 alone; it remains silent, by construction, on the positive case.

---

6. The Homological/Loop-Level Gap — the Open Problem

Here is the paper's central honesty, and the place a reader must not be allowed to slide past. The homological condition of §4 is **necessary but not sufficient** for refinement realisation. The two images are different objects:

- $\tau_*(H_1(\Gamma\_MS))$ is the image **in homology** — the set of classes hit by some 1-*chain* coming from Γ_MS.
- Im(τ_loops) is the image **at the loop level** — the set of classes hit by a single *reversible loop*.

Membership in the first does not imply membership in the second. A homology class can be represented as τ_* of a class without that class being representable by one reversible loop, for three distinct reasons, each a place the loop-level image can be strictly smaller:

1. **Chains versus loops.** $\tau_*(H_1)$ is generated by images of homology classes, which are combinations of cycles; realisation demands a *single* loop. A class reachable only as a sum of transported loops, with no single reversible loop mapping onto it, sits in $\tau_*(H_1)$ but outside Im(τ_loops).
2. **Reversibility restriction.** The domain is Loop_rev, reversible loops only. The homological image $\tau_*(H_1(\Gamma\_MS))$ is computed over all of H_1, which is generated without regard to which classes have reversible-loop representatives. Reversibility can exclude exactly the representatives that would realise γ_D — and this is the category tension the RC programme already isolated: γ_D is a *commitment* object, irreversible and downstream of refinement, while w must be a *pre-commitment reversible* loop. The reversibility constraint on the domain is precisely what makes "is γ_D loop-realised?" harder than "is [γ_D] in the homological image?"
3. **Torsion and boundaries.** Passage to homology quotients by boundaries and can collapse torsion distinctions that the loop-level question still sees; a class trivial or merged in H_1 may correspond to loop-level data that is not.

So the honest logical situation is:

realised $\Longrightarrow$ $[\gamma\_D] \in \tau_*(H_1)$ (Theorem 4.1, necessary), $[\gamma\_D] \in \tau_*(H_1) \not\Rightarrow$ realised (the gap, this section).

The residual difficulty of the entire programme lives in that second line. Closing it — turning the necessary condition into a sufficient one — is **loop-realisability of a homology class under the reversibility constraint**: given that $[\gamma\_D] \in \tau_*(H_1)$, decide whether a *single reversible* loop attains it. That is the open problem, and naming it is this paper's job. It is not closure topology (γ_D's persistence is settled), not phase theory, and not the homological computation of §4 (which gives only the necessary half); it is a property of the reversible loop space of Γ_MS under τ, and it is where all remaining uncertainty is now concentrated.

We therefore do **not** claim, as a looser treatment might, that "one needs only the image of τ." One needs the *loop-level* image; the *homological* image is a computable upper bound on it, useful for negative verdicts (Corollary 4.2) and necessary for positive ones, but not equal to the object realisation requires.

7. What Transfers, and What Does Not

The importance of refinement realisation is what would transfer with it. Here too the honest scope is narrower than "everything."

Theorem 7.1 (Transport-natural transfer) (proven)

If γ_D is refinement-realised by a reversible loop w, then every closure-side invariant that is **natural under transport** — every invariant I that factors through τ, so that its value is determined by the transported image — transfers to an invariant of admissible refinement motion, via pullback along w.

Proof. Let I be τ-natural: its value on a cycle depends only on that cycle's transported class, i.e. I is a function of $[\tau_{loops}(\cdot)]$ alone. Realisation gives a reversible loop w with $[\tau_{loops}(w)] = [\gamma_D]$; assign to the refinement motion the value $I(\gamma_D)$, pulled back along w.

The step that must be discharged, not assumed, is well-definedness: $[\gamma_D]$ may have several reversible-loop preimages, and the assigned value must not depend on which one is chosen. Suppose w and w′ are two reversible loops both realising γ_D, so $[\tau_{loops}(w)] = [\tau_{loops}(w')] = [\gamma_D]$. Because I depends only on the transported class, the pulled-back value along either loop is I evaluated on that common class, $[\gamma_D]$ — the same value for w and w′. So the assignment is independent of the choice of realising loop; it is a function of $[\gamma_D]$, not of the particular preimage. That independence is exactly what makes the pulled-back quantity an invariant of the refinement motion rather than an artefact of a chosen loop. (This is precisely where τ-naturality does its work: an invariant *not* factoring through τ — one sensitive to w beyond its transported class — would generally take different values on w and w′ and so would fail to be well-defined here; §7's "what does not transfer" is the converse of this same point.) Hence I transfers to a well-defined invariant of admissible refinement motion. ■

What does not transfer, and the qualifying instances

Theorem 7.1 deliberately does **not** say "every closure-side invariant." Pullback along τ transfers only what τ respects. An invariant sensitive to how γ_D sits inside C_{cl} in ways τ forgets — structure beyond the transported image — is not delivered by realisation, because τ does not carry it. The unqualified "every invariant" claim would prove more than pullback gives.

The invariants that *do* qualify are exactly the ones the programme cares about, which is why the narrowing costs nothing:

- **κ-charge** transfers: it is a transport-respected holonomy — κ is read through τ, the transport identification established in the companion phase paper **(inherited)**. We do not re-derive that identification here; we import it, and flag it as the one external dependency

of the transfer argument. Granting it, κ factors through transport and Theorem 7.1 applies.
- **Closure holonomy** transfers: holonomy is by construction a transport quantity, τ-natural by definition.
- **RC_path** survives: it is the loop-level route-dependence the residue *is*, so realisation makes it a refinement-level object.

A point of consistency with §3's choice of target must be stated, because §6 reason 3 (torsion/boundaries) might otherwise seem to reopen it. Every invariant in the list above is **class-level** — a function of the transported *homology class*, not of the particular cycle: κ and closure holonomy are homological by construction, and RC_path is route-dependence read at the level of loop classes. This is exactly why **class-level realisation suffices to transfer them**, and it closes the loop with §3, which adopted class-level realisation as primary. The torsion remark of §6 reason 3 is therefore describing the *gap* (homological image versus loop-level image), not arguing that the transferred invariants need the stronger on-the-nose target. None of them is sensitive to the cycle beyond its class; so the on-the-nose realisation §3 set aside is not required for any transfer the programme actually wants. Were some future invariant cycle-sensitive rather than class-sensitive, it would fall outside Theorem 7.1 (by the well-definedness argument above, which uses only the transported class) and would reopen the on-the-nose question — but no such invariant is among κ, holonomy, or RC_path.

Corollary 7.2 (What realisation would buy) (conditional on realisation)

If γ_D is refinement-realised then κ survives natively, closure holonomy becomes transport-visible, RC_path survives as a refinement-level residue, and the Memory Residue interpretation of phase becomes a refinement-level object rather than a closure-side-only one. Each of these is a τ-natural invariant, so each falls under Theorem 7.1; none requires the over-strong "every invariant" claim.

---

8. The Final Gate-3 Hinge

Three reductions converge on one point. The persistence papers reduced the closure question to γ_D. The Reversible-Connectedness papers reduced the native residue to transport. The phase paper reduced phase-as-memory to residue transfer. After the present paper, the remaining load-bearing question — stated at its correct, loop-level strength — is:

Is [γ_D] in the loop-level image of reversible admissible transport, Im(τ_loops)?

We state plainly what this paper has and has not shown about that question, because the "nothing else is load-bearing" temptation is exactly what the necessary-not-sufficient gap forbids.

- **Settled negatively-checkable (§4).** If $[\gamma\_D] \notin \tau_*(H_1(\Gamma\_MS))$, the answer is no, by a homology computation. The residue stays closure-side; no native transport residue survives; the Memory Residue interpretation fails as an account of substrate phase.
- **Open even given the homological condition (§5).** If $[\gamma\_D] \in \tau_*(H_1(\Gamma\_MS))$, the question is *not* thereby settled positively. Loop-level realisability under the reversibility constraint remains to be decided; this is the open problem.
- **Positive verdict, if achieved.** If a reversible loop w with $[\tau\_loops(w)] = [\gamma\_D]$ is exhibited (or its existence proved), then by Theorem 7.1 and Corollary 7.2 the τ-natural residue transfers: κ survives natively, RC_path survives, closure holonomy becomes transport-visible, and phase-as-memory becomes viable.

So "nothing else remains load-bearing" is true only in this precise sense: the one remaining load-bearing question is the loop-level image-containment, and it decomposes into a checkable necessary half (done) and an open sufficient half (named). Everything else — closure topology, persistence, the phase interpretation — is downstream of it.

---

9. What Must Be Computed

The problem is now completely localised, and the localisation is honest about its two stages.

Stage one (necessary, computable now). Compute the homological image $\tau_*(H_1(\Gamma\_MS)) \subseteq H_1(C\_cl)$ and locate $[\gamma\_D]$. If it lies outside, stop: realisation fails (Corollary 4.2), and the programme is settled negatively. This requires only homology — no new closure topology, no new persistence theory, no phase arguments.

Stage two (sufficient, open). If $[\gamma\_D]$ lies inside the homological image, the harder question begins: does a *single reversible* merge–split loop w attain $[\gamma\_D]$? Equivalently, is $[\gamma\_D]$ in $\mathrm{Im}(\tau\_loops)$, not merely in $\tau_*(H_1)$? This is loop-realisability under the reversibility constraint — the open problem of §6. It is a transport-theoretic question about the reversible loop space of Γ_MS, and it is where the reversibility restriction (the pre-commitment/commitment category gap) does its work.

The two stages should not be conflated, and the temptation to treat stage one as the whole problem is exactly the error this paper exists to forestall. Stage one can *refute* realisation; only stage two can *establish* it.

---

10. Structural Evidence for a Rank-One Transport Image

(structural evidence; conjectural — a different epistemic level from §§1–9)

The preceding sections establish the *criterion* and prove what is provable about it. This section is of a different character, and the header marks it so: it records not a theorem but a **convergence pattern** across several branches of the VERSF programme, all of which point toward the homological image $\tau_*(H_1(\Gamma\_MS))$ being small — plausibly rank one. None of what follows is proved here. The reader should hold §§1–9 (criterion, proved/open) and §10 (evidence, conjectural) at different levels, and the section is written to keep them apart rather than let the evidence borrow the standing of the theorems.

10.1 The domain-side rank bound — the strongest single point

One observation is structurally stronger than the others and belongs first, because it is a near-bound rather than an analogy. The K_N refinement framework computes a *rank-one* first homology:

$H_1(K\_N) \cong \mathbb{Z}\langle[C\_N]\rangle$ **(established, for K_N)**,

a single distinguished cycle class generated by the outer cycle C_N. The refinement framework does not naturally produce a large family of independent persistent cycles; it produces one topological persistence direction.

The consequence, *if* the K_N computation transfers to the transport domain Γ_MS, is immediate and strong — and it is an implication that needs no convergence argument. We state it formally, since it is the strongest intuition in the section and deserves the standing of a proposition rather than a remark.

Proposition 10.0 (Conditional rank bound) (proven, conditional on the hypothesis)

If $H_1(\Gamma\_MS) \cong H_1(K\_N)$, then rank $\tau_*(H_1(\Gamma\_MS)) \leq 1$.

Proof. $H_1(K\_N) \cong \mathbb{Z}\langle[C\_N]\rangle$ is rank one (established for K_N). By the hypothesis $H_1(\Gamma\_MS) \cong H_1(K\_N)$, so $H_1(\Gamma\_MS)$ is rank one. The image of a rank-one abelian group under any homomorphism has rank at most one, so rank $\tau_*(H_1(\Gamma\_MS)) \leq 1$. ■

The proposition is elementary, and deliberately so: its entire content is concentrated in its hypothesis. The implication itself is trivial group theory; what is not trivial is whether $H_1(\Gamma\_MS) \cong H_1(K\_N)$ — and that is precisely the thing the section flags as needing confirmation. Stating it as a proposition makes the logical structure exact: the rank-one upper bound is *one isomorphism away*, and that isomorphism is the single load-bearing input.

The bound is two-sided in consequence, and only one side is good news. This must be said plainly, because "≤ 1" reads as favourable when half of it is the worst possible outcome. Rank ≤

1 means the homological image is *either a single line or nothing* — rank 1 or rank 0 — and these are opposite verdicts for the programme, not a strong and a weak version of one verdict:

- **rank 1** is the programme-positive case: a single κ-line horizon, the necessary test sharpened to one projection, phase-as-memory still live.
- **rank 0** is the programme-*negative* case: the homological image is trivial, τ detects nothing, the horizon is empty, γ_D fails the obstruction for the most basic reason, and phase-as-memory is dead.

Discharging the K_N isomorphism therefore does **not** establish the favourable case. It establishes only that the answer is *one of these two* — a single line or nothing — collapsing the possibilities to a binary it does not resolve. Which of rank 0 or rank 1 holds is the non-triviality question, and it is exactly what the identity evidence of §§10.3–10.6 bears on: that evidence argues the line is *non-zero* (the persistent sectors exist and are one-dimensional, not absent), i.e. it argues for rank 1 over rank 0. So the two inputs do different jobs and neither substitutes for the other: Proposition 10.0's hypothesis (the K_N transfer) would prove rank ≤ 1, bounding the answer to {line, nothing}; the §§10.3–10.6 identity evidence argues which of the two, i.e. that the bound is attained rather than collapsing to zero. Only both together give rank exactly one.

This would settle the *upper half* of the rank-one picture outright: the homological horizon (§5, Proposition 5.4) would be at most a single line, regardless of any persistence-sector convergence. The remaining work would then be only to show the image is rank *exactly* one (non-zero) rather than rank zero — which is what §§10.3–10.6 bear on, and which the bound alone leaves undecided.

The load-bearing assumption is therefore the hypothesis of Proposition 10.0: **that the K_N computation transfers to Γ_MS**, i.e. H_1(Γ_MS) ≅ H_1(K_N), or more weakly that H_1 transfers between them. We do **not** assert K_N = Γ_MS here; if the isomorphism holds, the rank bound rank $\tau_*(H_1) \leq 1$ is proved (Proposition 10.0), with only non-triviality left to settle. If they are distinct frameworks whose homologies might differ, the K_N result is suggestive of the domain's shape but does not bound it, and the rank-one picture rests on the weaker convergence evidence below. Which case holds is itself a definite question — *does H_1(K_N) compute H_1(Γ_MS)?* — and it is the first thing to check, because a positive answer discharges the hypothesis of Proposition 10.0 and converts most of this section from evidence into a proved upper bound.

10.2 The H_1-Transfer Test

Proposition 10.0 makes the rank bound *one isomorphism away*; this section turns that isomorphism from an assumption into a testable condition, and decomposes the test into checkable parts. It is the operational unpacking of 10.0's hypothesis — what discharging it actually requires.

The desired identification is

$H_1(\Gamma\_MS) \cong H_1(K\_N) \cong \mathbb{Z}\langle[C\_N]\rangle$.

This cannot be assumed merely because both structures arise in the same refinement programme; it must be tested.

A comparison map is presupposed

The test is stated relative to a given comparison map $\iota : K\_N \to \Gamma\_MS$ — an inclusion if K_N is a subcomplex of the transport complex, or a constructed comparison otherwise. Whether such a map exists canonically is itself part of the transfer question: if K_N sits inside Γ_MS as a subcomplex the map is the inclusion and the test below is about its induced map on H_1; if the two are merely cousin objects of the refinement programme with no canonical map between them, then *constructing* an appropriate ι is a prerequisite to the test, not a given. We presuppose ι in what follows and flag its existence as the first thing the transfer question must settle, consistent with the rest of the paper's care that the K_N ↔ Γ_MS relationship is exactly the unverified link.

Definition (H_1-transfer)

K_N **H_1-transfers** to Γ_MS if the comparison map $\iota : K\_N \to \Gamma\_MS$ induces an isomorphism

$\iota_* : H_1(K\_N) \longrightarrow H_1(\Gamma\_MS)$.

Proposition 10.1 (H_1-transfer criterion) (proven)

K_N H_1-transfers to Γ_MS — i.e. $H_1(\Gamma\_MS) \cong \mathbb{Z}\langle[C\_N]\rangle$ — **if and only if both**:

(i) *(generation)* every reversible merge–split loop w in Γ_MS satisfies $[w] = n[C\_N]$ in $H_1(\Gamma\_MS)$ for some $n \in \mathbb{Z}$; and

(ii) *(freeness / non-torsion)* no non-zero multiple of [C_N] bounds in Γ_MS — equivalently [C_N] has infinite order in $H_1(\Gamma\_MS)$.

Proof. (⟸) Suppose (i) and (ii). By (i), [C_N] generates $H_1(\Gamma\_MS)$, so ι_* is surjective. By (ii), the only n with $n[C\_N] = 0$ is $n = 0$, so the map $\mathbb{Z} \to H_1(\Gamma\_MS)$, $n \mapsto n[C\_N]$, is injective; since $H_1(K\_N) = \mathbb{Z}\langle[C\_N]\rangle$, this is ι_*, so ι_* is injective. Hence ι_* is an isomorphism and $H_1(\Gamma\_MS) \cong \mathbb{Z}\langle[C\_N]\rangle$. (⟹) If ι_* is an isomorphism then $H_1(\Gamma\_MS) \cong \mathbb{Z}\langle[C\_N]\rangle$: every class is an integer multiple of [C_N] (giving (i)), and [C_N] generates a free $\mathbb{Z}$ (giving (ii)). ■

Why both conditions are needed — generation is not freeness. Condition (i) alone does *not* give the iff, and this is the one place the criterion must not be shortened to "every loop is a multiple of [C_N]." Generation is consistent with [C_N] having *finite* order: if some relation in Γ_MS forced $m[C\_N] = 0$ for a least $m > 1$, then every loop would still be a multiple of [C_N] — condition (i) would hold — yet $H_1(\Gamma\_MS) \cong \mathbb{Z}/m$, *not* $\mathbb{Z}\langle[C\_N]\rangle$. Generation gives surjectivity of ι_*; freeness (ii) gives injectivity; the iff needs both, and a treatment listing only (i) proves surjectivity and silently assumes injectivity. This is the same generation-versus-freeness distinction the paper draws elsewhere (the $\mathbb{R}$-vs-$\mathbb{Z}$ point of the appendix, the rank-≤1-includes-0 point of §10.1): "generated by [C_N]" is weaker than "freely generated by [C_N]," and only the latter is transfer.

There is a resonance worth flagging in the non-torsion condition. A torsion outcome m[C_N] = 0 is not the harmless failure of transfer; it is the appearance of *finite cyclic structure in the domain's own first homology* — and in a programme whose operative modulus is seven, the question "does [C_N] have infinite order, or order 7?" is not idle. A domain-side $\mathbb{Z}_7$ would hide exactly here, in the freeness check. So condition (ii) is not bookkeeping: it is where the programme's characteristic modulus could re-enter on the domain side, and it should be checked, not assumed.

Operational form of the test

The criterion separates into three checks; the first two establish generation (i), the third must additionally confirm freeness (ii).

1. **Generator coverage.** Every reversible merge–split move must be representable as motion along the K_N refinement skeleton. If Γ_MS contains a reversible move not representable in K_N, then Γ_MS may carry cycles invisible to K_N, and generation can fail.
2. **Relation completeness.** Every relation among reversible merge–split moves must be generated by the 2-cell relations already present in K_N. Extra relations in Γ_MS may collapse [C_N] (threatening freeness); missing relations may create extra homology (threatening generation).
3. **No extra cycles, and no collapse.** Every closed reversible merge–split word must reduce to a power of the primitive refinement cycle, $w \sim C\_N^n$ (generation) — *and* no non-trivial power C_N^m may itself reduce to the identity (freeness). The first clause is "H_1(Γ_MS) is generated by [C_N]"; the second is "[C_N] has infinite order." Both are required; check 3 fails transfer if either the word fails to reduce to a power of C_N, or a non-zero power of C_N reduces to nothing.

Consequence for the realisation criterion

If the H_1-transfer test succeeds — both (i) and (ii) — then Proposition 10.0 is active and rank $\tau_*(H_1(\Gamma\_MS)) \leq 1$. The homological image is therefore either

$\tau_*(H_1(\Gamma\_MS)) = 0$ or $\tau_*(H_1(\Gamma\_MS)) = \langle\kappa\rangle$,

the latter provided the transported generator is non-zero and identified with κ (the non-triviality and identity inputs, §§10.3–10.9 below). So the stage-one problem collapses to a binary: either admissible transport sees no native homological residue, or it sees exactly one residue line.

This binary carries the asymmetry the rest of the paper insists on, and the closing phrasing must not soften it: **"nothing" (rank 0) is the negative verdict for the whole programme** — empty horizon, γ_D unreachable, phase-as-memory dead — *not* a weaker form of the positive "one line." The H_1-transfer test, even when it succeeds, does not choose between the two; it only establishes that the answer is one of them. Choosing between them is the non-triviality question ($\kappa \neq 0$) that the persistence evidence below bears on.

What the test does not do

The H_1-transfer test does **not** prove refinement realisation. Even if $H_1(\Gamma\_MS) \cong H_1(K\_N)$, and even if $\tau_*(H_1(\Gamma\_MS)) = \langle\kappa\rangle$, one still needs the loop-level result: a reversible merge–split loop w with $[\tau\_loops(w)] = [\gamma\_D]$. H_1-transfer closes the *rank* problem (reducing the homological horizon to "line or nothing"); it does not close the *loop-realisation* problem. Its value is precisely and only that reduction — the Stage-One/Stage-Two boundary is untouched, as everywhere in this paper.

10.3 The persistent direction of the σ-sector

The continuum-limit analysis of sequential transport establishes that the σ-sector possesses a unique persistent mode. In original spoke variables this is the alternating mode; under the carrier–envelope decomposition it becomes the constant-envelope mode, which forms the kernel of the continuum transport operator. The continuum theory therefore contains a distinguished one-dimensional persistence sector, surviving the continuum limit and not removed by admissibility-restoring flow. **(established, for the σ-sector)**

This is a rank-one persistent structure — but on its own it is evidence that *the σ-sector* is rank one, not yet that its persistence class is the *same* class as the transport image. That identification is part of the conjecture, not an input to it.

10.4 The persistent cohomological sector

Independent transport papers identify a refinement-persistent cohomological sector represented by the κ-class, which survives refinement and carries transport residue. The σ-sector paper itself flags the relation between the continuum σ-sector and this persistent cohomological transport sector as an **explicitly open question**.

We state the weight of that flag carefully, because it is easy to overread. That two independent constructions each find a persistent, one-dimensional sector, and that the programme has *posed* their relation as a question worth asking, is a reason to investigate whether they are identical — not evidence that they are. "The relation is an open question two papers thought worth raising" is an invitation to the conjecture below, not partial confirmation of it. What is genuinely true is the weaker statement: the σ-sector and the κ-sector are independently identified as persistent and one-dimensional, and their identification is a posed open problem.

10.5 Closure residue and homotopy persistence

Independent closure-phase work establishes that closure invariants factor through homotopy classes of admissible cycles and remain non-trivial whenever the homotopy class does. This gives a second route by which persistent transport information attaches to persistent cycle classes. The resulting picture is that transport persistence, closure persistence, and homotopy persistence all *appear to* track a one-dimensional persistent object.

10.6 The single global closure mode

A further piece of structural evidence comes from the K = 7 closure-symmetry analysis, and it points the same way without being overstated. That analysis decomposes the closure channels as

11 boundary channels + 2 vertex channels + 1 global mode = 14,

and the figure 14 is not the point. The point is the last summand: there is exactly **one global closure mode**. The boundary and vertex channels are local; only a single mode is global. If the transport-visible residue is carried by global structure — which is the whole thrust of the programme (Proposition 5.4's horizon is a statement about global homology) — then a closure architecture with one global mode is exactly the kind of structure whose surviving transport image would be rank one. A single global closure mode is the channel-level shadow of a rank-one global residue.

We keep this at evidence-strength and do not lean on it harder than it bears: that there is one global closure channel is consistent with, and suggestive of, a rank-one transport image; it is not a computation of that image, and it does not by itself identify the global mode with κ or with the σ-kernel. It is a third independent appearance of "exactly one" — one global channel, alongside the one refinement cycle (§10.1) and the one persistent σ-mode (§10.3) — and its value is precisely that the "one" recurs across independent decompositions.

10.7 The inference, and its limit

Here the section must be most careful, because the cumulative rhetoric of §§10.1 and 10.3–10.6 builds a momentum the evidence does not yet license. Four branches each exhibiting a rank-one persistent sector is **consistent with** their being one shared class — but it is equally consistent with four *distinct* rank-one classes that happen to coexist. Co-occurrence of four one-dimensional sectors is not evidence of their identity; identity is a strictly stronger claim, and it is exactly what the conjecture below asserts and what remains unproven.

So the honest reading of §§10.1 and 10.3–10.6 is: the *upper bound* (rank ≤ 1) is near-proved from the domain side *if* the K_N transfer holds (§10.1), while the *identity* of the four persistent sectors — the claim that they are one class, hence that the image is non-zero and is the κ-line — is convergent evidence only, not yet established.

10.8 Conjecture (Persistent-Class Identification)

Conjecture 10.2 (Persistent-Class Identification) (conjectural)

The persistent σ-kernel class, the refinement-persistent cohomological class κ, and the transport-visible closure residue are different descriptions of a single underlying persistence class. Equivalently, up to canonical identification,

$\langle \kappa \rangle = \tau_*(H_1(\Gamma\_MS))$

(an equality of *homological* images). If true, the homological transport image is rank one, generated by κ.

What rank-one would and would not settle. This must be stated against the §6 gap, or it reads as nearly solving the whole problem when it sharpens only one half. If Conjecture 10.2 holds, the homological obstruction of §4 collapses to a single projection: the test "$[\gamma_D] \in \tau_*(H_1)$?" becomes "is $[\gamma_D]$ a non-zero multiple of κ?" — one scalar question on one line. That is a dramatic sharpening of **stage one** (the homological, necessary half). It leaves **stage two entirely untouched**: even granting $\tau_*(H_1) = \langle\kappa\rangle$ rank one, whether γ_D is *loop-realised* within that line — whether a single reversible loop attains it, not merely whether its class lies on the κ-line — is the §6 loop-level problem, open exactly as before. Rank-one would tell us the horizon is a single line and make the necessary test trivial to state; it would not tell us whether γ_D is reached within it. The closing phrase "γ_D projects non-trivially onto the persistence direction" is therefore the *homological* test (Theorem 4.1 again), not loop-realisation — sharpened, but still only the necessary half.

10.9 Status

No proof of Conjecture 10.2 is supplied. The paper establishes neither equality nor isomorphism between the σ-sector persistence class and κ. What is established is the convergence pattern:

1. the K_N refinement framework has rank-one first homology **(established for K_N; transfers to Γ_{MS} only if the domains agree, §10.1)** — the strongest point, a near-bound on the image rank from the domain side;
2. the σ-sector contains a unique persistent mode (established for the σ-sector);
3. the transport sector contains a unique persistent cohomology class κ;
4. closure persistence is controlled by homotopy persistence;
5. the $K = 7$ closure architecture contains exactly one global mode (§10.6) — corroborating structure rather than a fifth independent sector, but a further independent appearance of "exactly one".

The conjunction suggests the transport-visible residue may be substantially smaller than the full closure homology — plausibly cyclic, possibly rank one. But the two honest limits stand: the upper bound is a near-theorem only if the $K_N \to \Gamma_{MS}$ transfer holds (otherwise it is evidence), and the identity of the persistent sectors is convergent evidence, not proof — four one-dimensional sectors might still be four different one-dimensional sectors. Determining whether the identification is correct, and whether the K_N homology computes the transport domain's, remain open — the first being the cleaner of the two to attempt, since it could convert §10.1 from evidence (Proposition 10.0's hypothesis unverified) into a proved rank bound (its hypothesis discharged).

10.10 Conditional closure of the homological image

It is worth stating the strongest closure presently available as a single theorem, because it bundles the section's scattered conditional inputs into one explicit statement and makes exact what discharging them would — and would not — buy. The theorem introduces no new assumptions: its four hypotheses are precisely the inputs §§10.1–10.9 have already isolated, collected in one place.

Theorem 10.3 (Conditional homological closure) (proven, conditional on (1)–(4))

Assume:

1. $H_1(\Gamma\_MS) \cong H_1(K\_N)$ — the K_N transfer (hypothesis of Proposition 10.0, §10.1);
2. $H_1(K\_N) \cong \mathbb{Z}\langle[C\_N]\rangle$ — the K_N rank-one computation (established for K_N, §10.1);
3. $\tau_*([C\_N]) = \kappa$ — the transported generator is the κ-class (the identity content of §§10.3–10.8);
4. $\kappa \neq 0$ — non-triviality (the κ-line is not the zero line, the rank-1-not-rank-0 selection of §10.9).

Then $\tau_*(H_1(\Gamma\_MS)) = \langle\kappa\rangle$.

Proof. By (1) and (2), $H_1(\Gamma\_MS)$ is rank one; let a be a generator corresponding to [C_N] under the isomorphism of (1). Every class $x \in H_1(\Gamma\_MS)$ is then $x = na$ for some $n \in \mathbb{Z}$, so $\tau_*(x) = \tau_*(na) = n\,\tau_*(a)$. By (3), $\tau_*(a) = \kappa$, hence $\tau_*(x) = n\kappa$: every element of the image lies in $\langle\kappa\rangle$. Conversely every $n\kappa$ is $\tau_*(na)$, so $\langle\kappa\rangle$ lies in the image. Thus $\tau_*(H_1(\Gamma\_MS)) = \langle\kappa\rangle$, and by (4) this is non-trivial and rank one. ■

How the four hypotheses map to the section's two open inputs

The theorem is not a new assumption set; it is the conjunction of the two inputs §10 has discussed throughout, split into checkable pieces. Hypotheses (1)–(2) are the **transfer** — Proposition 10.0's hypothesis plus the K_N computation — which alone gives only rank ≤ 1, i.e. "line or nothing" (§10.9). Hypotheses (3)–(4) are the **identity-and-non-triviality** content — that the one available direction is specifically the κ-line and that this line is non-zero — which is exactly what §§10.3–10.8 argue for as convergent evidence and Conjecture 10.2 asserts. So Theorem 10.3 is the precise statement of "what both inputs together would give," with (1)–(2) bounding and (3)–(4) selecting, matching the two-input structure §10.9 set out. None of the four is proved here; the theorem states their joint consequence.

Corollary 10.4 (Reduced homological test) (proven, conditional on (1)–(4))

Under the assumptions of Theorem 10.3 the homological obstruction reduces to a single κ-line test:

$$[\gamma\_D] \in \tau_*(H_1(\Gamma\_MS)) \Leftrightarrow [\gamma\_D] \in \langle\kappa\rangle \Leftrightarrow \exists\, n \in \mathbb{Z} \text{ with } [\gamma\_D] = n\kappa.$$

In the $\mathbb{Z}_7$ sector this is $[\gamma\_D] = n\kappa \pmod 7$ for some $n \in \mathbb{Z}_7$. If additionally κ *generates* the $\mathbb{Z}_7$ sector — a fifth condition, not among (1)–(4), and to be checked separately — the test is equivalent to [γ_D] carrying non-zero κ-charge.

Status — Stage One only

Theorem 10.3 does **not** prove refinement realisation. It conditionally closes the *homological image*, and the Stage-One/Stage-Two boundary the whole paper turns on is preserved exactly. What it would establish, given (1)–(4), is that the homological horizon of admissible transport (§5.4) is *exactly* the κ-line — neither larger (the upper bound, from (1)–(2)) nor zero (non-triviality, from (4)). What it leaves entirely untouched is the loop-level question:

does there exist a single reversible merge–split loop w with [τ_loops(w)] = [γ_D]?

That is unchanged by the theorem. Even with the horizon closed to exactly ⟨κ⟩ and [γ_D] confirmed to lie on it, whether a reversible loop *attains* [γ_D] — rather than merely [γ_D] sitting in the line that loops could in principle reach — is the §6 gap, open as before. So Theorem 10.3 would close Stage One (the necessary, homological half) completely: the horizon is the κ-line, and the obstruction test is a single κ-charge check. It would not advance Stage Two (the sufficient, loop-level half) at all. The paper's central asymmetry holds to the end: closing the homological image bounds and sharpens the necessary condition without touching the sufficient one.

---

11. Conclusion

The native-residue question has been reduced to a single transport criterion and that criterion has been correctly scoped. Refinement realisation is, by definition, the membership of [γ_D] in the loop-level image Im(τ_loops). That restatement is not itself a result; the results are two and they bracket the open problem from both sides.

From above: a **homological obstruction** (Theorem 4.1) gives a necessary condition, [γ_D] ∈ $\tau_*(H_1)$, checkable by computation, and a disqualifier (Corollary 4.2) that can settle the question negatively without any loop-space analysis. From the consequence side: a **transport-natural transfer** theorem (Theorem 7.1) establishes that realisation, if it holds, carries exactly the τ-natural invariants — κ, holonomy, RC_path — and no more.

Between them lies the open problem, named here as such: the **gap between the homological image and the loop-level image**, i.e. loop-realisability of [γ_D] under the reversibility constraint. Membership in homology is necessary; membership achievable by a single reversible loop is what realisation requires; the two are not the same, and the difference is the reversibility restriction that encodes the commitment/pre-commitment category gap.

The closure-side existence and persistence of γ_D are settled and separate. What remains is neither where a native residue might live nor whether γ_D survives, but whether reversible admissible transport *reaches* the cycle whose persistence is already secured. This paper does not answer that. It proves the half that is provable, names the half that is open, and leaves the programme with one concrete computation to attempt first (the homological image) and one concrete theorem to seek after (reversible loop-realisability) — in that order, because the first can end the question and the second is needed only if it does not.

A final word on direction, kept deliberately at evidence-strength. §10 records a convergence across the programme suggesting the homological image may be rank one — generated by the single κ-class — with the strongest single indication being the rank-one first homology of the refinement framework, which would bound the image rank directly if it transfers to the transport domain (Proposition 10.0). That bound, however, is two-sided and only half of it is favourable: it would force the image to be *a single line or nothing*, rank 1 or rank 0, and the K_N transfer alone does not choose between them — rank 0 is the programme-negative outcome (empty horizon, γ_D unreachable, phase-as-memory dead), rank 1 the positive one. The identity evidence of §10 argues for the non-zero case; it does not prove it. So even the favourable reading rests on two separate steps — the transfer (bounding to {line, nothing}) and non-triviality (selecting the line) — collected as the four hypotheses of Theorem 10.3 (§10.10), whose joint consequence would be that the horizon is *exactly* the κ-line. Were all four to hold, the *necessary* test would collapse to a single projection onto that line, sharpening the question's first half while leaving the second untouched. Whether the image is rank one rather than rank zero, and whether γ_D is loop-realised within it if so, are the open problems beyond the criterion this paper sets. The criterion is proved; the answer it points toward is, for now, only glimpsed — and glimpsed includes the possibility that the horizon is empty.