Conditional Closure of Gate-3 Transport

The Assumption Set, the Homological Verdict, and the Single Remaining Hinge

Keith Taylor — VERSF Theoretical Physics Programme

General Reader Summary

What survives, and on exactly which promises.

A sequence of papers reduced the question of whether reality keeps a permanent trace of irreversible events to a short list of precise conditions. This paper does the honest closing thing: it states plainly what follows if *those conditions hold, lists each condition with where it stands and what would settle it, and stops — without pretending the conditions are proved, and without circling them endlessly.*

Over several papers the question "does reality keep a permanent topological memory of irreversible events?" was sharpened from a vague idea into a chain of precise conditions. Each paper either established a link or reduced it to something more concrete. At this point the chain is short enough to lay out in full, and the responsible thing is to do exactly that: say what the chain *delivers* if its remaining links hold, name each link honestly, and move on to the next layer of the programme rather than spend further papers rephrasing the same unproven step.

This paper collects four assumptions, calls them A, B, C, and D, and proves a single clean conclusion from them: if all four hold, then the "memory" the programme has been chasing survives in the precise homological sense — it lives on exactly one line, the κ-line, and several downstream structures (including the phase-as-memory interpretation of quantum phase) remain viable. That is the positive verdict.

The honesty is in the bookkeeping. Each of the four assumptions is given with its current status: which earlier paper established or reduced it, what specifically would have to be done to *settle* it, and — crucially — what happens to the verdict if it turns out false. Two of the four are close to secured (one is proved outright under a stated locality condition; another is reduced to a single combinatorial property with a known strategy). Two are genuine open inputs about how transport reads the loop. None is dressed up as more settled than it is.

But the paper also shows the open pieces are far from arbitrary, and this is its main new contribution. Every structure the whole programme has ever produced points in a *single* direction — one loop, one mode, one class, one global channel — and no construction has ever turned up a second, independent one. So the live question is not "is the answer one of infinitely many shapes?" but the much narrower "is it the single shape everything points to, or that same shape plus some extra piece nobody has ever seen?" The burden flips: the simple answer is now the default, and any more complicated answer is the one that would have to produce a witness. Sharper still, the two genuinely open conditions — that the loop's image lands on the expected line, and that it does not vanish — turn out, once one grants that the relevant space is a single line, to be *the same question* wearing two hats: does the loop project onto that line with a non-zero amount, or with zero? Non-zero, and the memory survives; zero, and it collapses. So the two open inputs are really one number, and the question is only whether that number is zero.

And one thing is deliberately kept *outside* the assumption list, because it is not one of the conditions the closure rests on — it is the thing the closure leaves open. Even if all the assumptions hold and the memory survives on the expected line, there remains a further, harder question: whether a single reversible motion actually *reaches* the specific trace γ_D, rather than merely the trace lying on the line that motions could in principle reach. That loop-level question is the final hinge of the whole programme, and this paper does not touch it. The closure here is homological — it settles where the memory lives and that it is one line — not the loop-level realisation that would settle whether the memory is actually attained.

So the structure is: here is the conditional closure; here are the assumptions, each with its status and its escape route; here is how tightly the open ones are already pinned — one simple shape that everything points to, with the last genuinely-open question reduced to whether a single number is zero; and here is the one question that remains beyond all of them. If future work confirms the assumptions, the conclusion follows. That is an endpoint for this branch — not a completed proof, and not an endless recursion on the same step, but an honest statement of what has been reduced to what, and of how little room the remaining unknowns now have.

Abstract

This paper states the conditional closure of the Gate-3 transport branch. It does not prove the closure unconditionally; it assembles the results and reductions of the preceding papers into a single explicit assumption set A–D, proves the homological verdict that follows from it, and maps each assumption to its prior-paper status, its discharge condition, and the consequence of its failure.

The assumptions are: **(A)** locality of the reversible merge–split calculus, giving $\mathrm{ord}(C\_N) = \infty$ (the freeness result, *Order of the Primitive Refinement Cycle*); **(B)** the H_1-transfer condition $H_1(\Gamma\_MS) \cong \mathbb{Z}\langle[C\_N]\rangle$, equivalently that every reversible loop reduces to a power of C_N (the generation reduction, *Generation Theorem for Reversible Merge–Split Transport*); **(C)** identification of the transported generator, $\tau*([C\_N]) = \kappa$; **(D)** non-triviality, $\kappa \neq 0$. Under A–D we prove $\tau*(H_1(\Gamma\_MS)) = \langle\kappa\rangle$: the homological horizon of admissible transport is exactly the κ-line. The Gate-3 verdict follows — the native residue survives homologically, κ survives, RC_path survives, and the phase-as-memory interpretation remains viable.

Two disciplines are preserved. First, the verdict is **two-sided**: D is not a strength dial but a switch — D failing ($\kappa = 0$) is the programme-*negative* outcome (empty horizon), not a weaker positive one. Second, loop-realisation of γ_D is **not among A–D**: it is not a condition the closure rests on but the single question the closure explicitly leaves open. Closing A–D settles the *homological* horizon (Stage One); whether a single reversible loop *realises* γ_D (Stage Two) is untouched and remains the final Gate-3 hinge. The paper's purpose is to mark this branch's endpoint honestly — conditional closure with an explicit, status-annotated assumption set — so the programme can advance to its next layer without recursively rephrasing the same unresolved transport step.

Epistemic markers: ***(established/proven)*** *for results secured in prior papers;* ***(reduced)*** *for a condition reduced to a concrete sub-problem;* ***(open input)*** *for an assumption not yet established;* ***(conditional)*** *for the present paper's verdict;* ***(out of scope)*** *for the loop-level question deliberately left open.*

Table of Contents

---

1. Purpose: Why a Conditional Closure

The Gate-3 transport branch has reached the point where its remaining content is a short list of precise conditions, each either established, reduced to a concrete sub-problem, or isolated as an explicit open input. When a branch reaches that state, the intellectually honest move is not to continue producing papers that rephrase the same unproven step, but to state the conditional closure plainly: *here is what follows if the conditions hold; here are the conditions, each with its status; if future work confirms them, the conclusion follows.*

This is acceptable here for a specific reason: the conditionality is already explicit throughout the branch. No careful reader of the preceding papers could mistake "if A–D hold, Gate-3 survives" for "A–D are proved" — each paper marked exactly where its conditionality lived, and this paper inherits and consolidates those markings rather than obscuring them. Stating the conditional closure therefore pretends nothing. What it avoids is the opposite failure: spending further papers recursively rephrasing "prove every reversible merge–split loop reduces to C_N^n" without gaining new information. At some point the honest thing is to record the conditional closure, list the assumptions, and move to the next layer of the programme.

This paper is that record. It is deliberately short and adds no new mathematics; its content is organisational — the assembly of prior results into one assumption set, the single verdict that follows, and an honest status map of what remains. Its value is as a branch endpoint and a frontier map, not as a new theorem.

---

2. The Assumption Set: A–D and the Reformulation A, B, O, Λ

The closure rests on exactly four assumptions, stated here in full and annotated in §5 with their status.

Assumption A (Locality → freeness). The reversible merge–split calculus is local: every admissible reversible relator is a local contractible refinement relation. Consequence (proved in the companion freeness paper via the constructed winding homomorphism): the primitive refinement cycle has infinite order,

$\mathrm{ord}(C\_N) = \infty$.

Assumption B (H_1-transfer / generation). The H_1-transfer condition holds:

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

equivalently, every reversible merge–split loop reduces to a power of C_N (generation), so [C_N] generates the domain homology and — with A giving freeness — generates it freely.

Assumption C (Generator identification). The transported generator is the κ-class:

$\tau*([C\_N]) = \kappa$.

Assumption D (Non-triviality). The transported residue is non-zero:

$\kappa \neq 0$.

A and B are domain-side (about $H_1(\Gamma\_MS)$ itself); C and D are transport-side (about how τ∗ carries the domain generator into the target). A is proved under its stated locality hypothesis; B is reduced to a single combinatorial property; C and D are open inputs. None is asserted here as unconditionally established — the verdict of §3 is explicitly conditional on all four.

2.1 The reformulated set: A, B, O, Λ

A–D is the set as the prior papers framed it, and we retain it because the compression below is stated in its terms. But §6.4 shows that C and D are not independent: once the transport-visible persistent sector is one-dimensional, the transported generator has the form $\tau_*([C_N]) = \lambda\kappa$, and C and D collapse into the single condition $\lambda \neq 0$. The genuine structural content of the closure is therefore carried by a cleaner four-element set, which we name and adopt as canonical:

- **A (Locality).** As above — gives $\mathrm{ord}(C_N) = \infty$, freeness.
- **B (H_1-transfer).** As above — $H_1(\Gamma_{MS}) \cong \mathbb{Z}\langle[C_N]\rangle$, generation + freeness.
- **O (One-dimensionality).** The transport-visible persistent sector is one-dimensional, spanned by κ.
- **Λ (Non-zero projection).** The primitive cycle projects non-trivially onto that sector: $\lambda \neq 0$, where $\tau_*([C_N]) = \lambda\kappa$.

The relation between the two presentations is exact and proved in §6.4: **(C and D) ⟺ (O and Λ)**, in the sense that O carries C's identification content (it is what fixes the line as κ's rather than some other cyclic direction — see §6.2) and Λ carries D's non-triviality, while Λ also delivers C's *direction* given O. So the closure may be read either way:

original: A, B, C, D canonical: A, B, O, Λ

with A, B domain-side and O, Λ transport-side. The canonical form is cleaner because its transport side is *one geometric hypothesis (O) plus one scalar bit (Λ)* rather than two entangled identification/non-triviality assumptions, and because it makes the genuine open content legible: O is a structural claim about the target sector (one-dimensional), Λ is a single non-zero check. We use A–D where continuity with the prior papers matters (the status map of §5, the discharge programme of §8) and the canonical A/B/O/Λ where the cleanest statement of what the branch rests on is wanted (the present subsection, the synthesis of §6.5). They are the same closure; the verdict of §3 holds under either presentation, since each implies the other.

---

3. The Homological Verdict

Theorem 3.1 (Homological horizon under A–D) (conditional on A–D)

Assume A, B, C, D. Then

$\tau_*(H_1(\Gamma_{MS})) = \langle\kappa\rangle$,

i.e. the homological horizon of admissible transport is exactly the κ-line.

Proof. By B, $H_1(\Gamma\_MS) \cong \mathbb{Z}\langle[C\_N]\rangle$: every class is n[C_N] for a unique $n \in \mathbb{Z}$ (uniqueness from A's freeness). Apply τ_*: every element of the image is $\tau_*(n[C\_N]) = n\cdot\tau_*([C\_N]) = n\kappa$ by C. Hence $\tau_*(H_1(\Gamma\_MS)) \subseteq \langle\kappa\rangle$. Conversely every $n\kappa = \tau_*(n[C\_N])$ lies in the image, so $\langle\kappa\rangle \subseteq \tau_*(H_1(\Gamma\_MS))$. Thus $\tau_*(H_1(\Gamma\_MS)) = \langle\kappa\rangle$, and by D this line is non-trivial. ■

This is the rank-one image the whole branch has been driving toward, now assembled: A and B fix the domain as a single free generator, C identifies its transported image as κ, and D makes that image non-zero. The horizon is neither larger than one line (A, B bound it) nor zero (D excludes collapse). It is exactly ⟨κ⟩.

---

4. The Gate-3 Verdict

Corollary 4.1 (Gate-3 homological survival under A–D) (conditional on A–D)

Under A–D, the following hold *at the homological level* — each is a statement about the surviving horizon ⟨κ⟩, not about realisation into refinement motion (which is Stage Two, §7, and outside A–D):

- the **native residue survives homologically** — the transport image is the non-zero line ⟨κ⟩, so admissible transport detects a genuine homological residue;
- **κ survives as the horizon generator** — the closure charge is the generator of the non-zero image line ⟨κ⟩;
- **RC_path has a surviving home** — the reversible-connectedness route-dependence is a τ-natural invariant, and under A–D there is now a non-zero ⟨κ⟩ for it to be supported on; whether it *transfers into refinement motion* is the Stage-Two realisation question, deferred, not concluded here;
- **phase-as-memory retains its substrate object** — the companion phase paper's interpretation requires a surviving transport residue, and ⟨κ⟩ is one; its viability at the homological level is secured, its realisation is not.

Proof. Each follows from Theorem 3.1. The residue and κ survive because the image is the non-zero ⟨κ⟩. RC_path and the phase interpretation are τ-natural invariants, and Theorem 3.1 supplies a non-zero generator ⟨κ⟩ for them to be supported on — which is what "survives homologically" means and all that A–D deliver. We do **not** invoke the criterion paper's transport-natural *transfer* theorem here: its hypothesis is realisation (Stage Two), which A–D do not establish (§7). The corollary's content is therefore that these invariants have a non-zero homological support under A–D, not that they transfer into refinement motion. ■

This is the positive Gate-3 verdict at the level A–D actually reach: the homological one. Under A–D the native-residue programme survives *homologically*, and the structures built on it — κ, RC_path, phase-as-memory — retain a non-zero ⟨κ⟩ to stand on. Whether they transfer into refinement motion is the Stage-Two realisation of γ_D, which §7 holds outside the assumption set; the firewall of §§6–7 applies here too, and §4 claims only the Stage-One half.

---

5. Status of Each Assumption

The value of a conditional closure is an honest status map. Each assumption is annotated with: its content, where it was established or reduced, what would discharge it, and what its failure would mean.

A — Locality → freeness. Status: (established/proven) under its locality hypothesis. *Established where:* the freeness paper, by constructing a $\mathbb{Z}$-valued winding homomorphism W with W([C_N]) = 1 and showing it is well-defined precisely when every relator is winding-zero, which locality guarantees. *Discharge:* confirm that the reversible merge–split relation set R_MS contains no global period relation (no relator of non-zero winding) — a definite question about the calculus's definition. *If it fails:* a global winding-m relator gives ord(C_N) = m (m = 7 the natural case), so $H_1(\Gamma$_MS) is finite cyclic and the seven sits in the *domain* rather than the readout. The closure as stated would not hold; a different (domain-seven) closure would. One cross-paper reconciliation is worth stating here, since the programme contains more than one seven and they must not be silently conflated: a domain-seven (ord(C_N) = 7) would be a *new* appearance of seven-fold structure, in the reversible merge–split domain, and is **not** the same object as the phase paper's $\mathbb{Z}_7$ transport-*residue* seven (which lives in the target, κ-valued, read after transport). If A fails, the question becomes whether the domain-seven and the target-residue-seven are the same K = 7 architecture seen at two stages (the natural expectation, given the programme's single K = 7 source) or two independent sevens (which would be structurally surprising and would itself need explanation) — exactly the domain-seven-versus-target-seven dichotomy the freeness paper raised. The present closure assumes A (locality), under which the domain is free and the only seven is the target residue; the A-failure branch is where this reconciliation would have to be carried out, and it is flagged, not resolved, here. This is the most nearly secured assumption: it is proved outright given locality, and only locality remains to be confirmed.

B — H_1-transfer / generation. Status: (reduced) to a single combinatorial property. *Reduced where:* the generation paper, which showed generation follows from termination (T), the closed-normal-form property (NF), and the bridge hypothesis (C-complete) that R_MS computes $H_1(\Gamma$_MS); and which identified (NF) — every closed irreducible is a C_N-power — as the real content, equivalent to "no extra independent cycles." *Discharge:* prove (NF) via the §11 measure strategy of that paper, with the winding-collection reduction step the located difficulty; and confirm the bridge hypothesis that the presented calculus is Γ_MS. *If it fails:* an extra closed irreducible is an extra free generator, so $H_1(\Gamma$_MS) ≅ $\mathbb{Z}^r$ with r > 1; the domain-side rank exceeds one and the rank-one image picture breaks from the domain side. This is the genuine domain-side open risk: torsion is excluded (A), extra rank is not (B).

C — Generator identification $\tau*([C\_N]) = \kappa$. Status: (open input). *Where it stands:* argued for as convergent evidence in the criterion paper's structural-evidence section (the σ-sector, the persistent cohomological κ-class, closure/homotopy persistence, the single global closure mode all pointing at one persistence direction), and asserted by that paper's Persistent-Class Identification conjecture — but not proved. *Discharge:* prove the persistent-class identification, i.e. that the transported domain generator is the κ-class rather than some other transport-visible class. *If it fails:* the horizon is still rank ≤ 1 (given A, B) but its generator is not κ; downstream structures keyed to κ specifically (the phase paper's residue identification) would need re-keying to whatever the true generator is. The rank-one picture survives; the κ-labelling does not.

D — Non-triviality $\kappa \neq 0$. Status: (open input), and the switch the verdict turns on. *Where it stands:* the convergence evidence argues the persistent direction is non-zero (the sectors exist, are one-dimensional, not absent), but non-triviality is not proved. *Discharge:* show the transported generator is genuinely non-zero — equivalently, that the rank-one bound is *attained* (rank exactly 1) rather than collapsing to rank 0. *If it fails ($\kappa = 0$):* this is the programme-negative outcome. The horizon is empty, τ detects no native residue, γ_D is homologically unreachable, and phase-as-memory is dead. D is not a strength parameter; it is the binary on which the entire positive verdict depends.

The map in one line: **A proved-given-locality; B reduced to (NF); C and D open inputs, with D the switch.** Torsion excluded, extra rank the domain-side risk, κ-identity and κ-non-triviality the transport-side inputs.

---

6. How Constrained the Open Pieces Already Are

The status map of §5 could be read as "two proved-or-reduced assumptions and two open ones," as though B's residual gap and C's identification were arbitrary hypotheses awaiting unconstrained verification. They are not, and saying why turns the closure from four loosely-related assumptions into one emerging picture whose weakest pieces are already heavily constrained. This section makes that precise — without proving B or C, which it does not, but by shifting the *default* against which each is judged.

6.1 B is the minimal-rank hypothesis

Proposition 6.1 (No evidence for additional generators) (structural; an argument from the current state of construction, not a proof of absence)

Every persistent structure that has appeared anywhere in the programme to date is rank one: the K_N homology ($H_1(K\_N) \cong \mathbb{Z}\langle[C\_N]\rangle$, established), the σ persistence mode (a unique continuum kernel), the κ persistence class (a single refinement-persistent cohomology generator), and the global closure mode (exactly one, in the $K = 7$ channel decomposition $11 + 2 + 1$). No construction in the programme has produced an independent second candidate transport generator.

What this does and does not establish. It does **not** prove B: absence of a known second generator is an argument from the present state of the programme's constructions, not a theorem that none exists, and a future construction could exhibit one. What it does is **shift the default.** Without it, a reader weighing B faces an open menu — rank 1, rank 2, rank 3, rank r for any r — with no reason to privilege one. With it, the menu collapses to a binary: *rank one*, or *rank one plus a currently-unknown extra generator that no construction has yet produced*. Those are very different epistemic positions. The burden moves: rank one is the natural expectation, and rank > 1 is the claim that now requires a witness — an explicit second generator — rather than rank one being the claim that must defend itself against an open field of alternatives. B's residual gap (the (NF) property of the generation paper) is thereby not "is the rank one of infinitely many possibilities" but "has the single missing generator, which would have to be exhibited, in fact failed to appear because it does not exist." That is a constrained gap, not an open one.

6.2 C and D compress to one-dimensionality plus a coefficient

This is the larger gain, because it is a genuine conditional implication rather than an argument from absence.

Proposition 6.2 (Uniqueness pressure on the transported class) (proven, conditional on H_1-transfer and one-dimensionality of the transport-visible sector)

Assume (B) H_1-transfer, so $H_1(\Gamma\_MS) \cong \mathbb{Z}\langle[C\_N]\rangle$, and assume the transport-visible persistent sector is one-dimensional, spanned by κ. Then the transported generator is proportional to κ:

$\tau_*([C\_N]) = \lambda\kappa$ for some coefficient λ.

Proof. By B, $H_1(\Gamma\_MS)$ is generated by [C_N], so its image under τ_* is generated by the single element $\tau_*([C\_N])$; the image is therefore cyclic. By hypothesis the transport-visible persistent sector — the part of the target transport can populate — is one-dimensional, spanned by κ. A cyclic subgroup of a one-dimensional sector spanned by κ lies in ⟨κ⟩; hence $\tau_*([C\_N])$ is a multiple of κ, $\tau_*([C\_N]) = \lambda\kappa$. ■

What this changes. The identification problem C is no longer "determine an arbitrary transport image $\tau_*([C\_N])$ out of all possible target classes." It is the far smaller problem **determine the coefficient λ** — a single scalar — given that the image is already forced onto the κ-line by rank-one-ness. And the two transport-side assumptions fuse: if κ generates the persistence sector and $\lambda \neq 0$, then $\tau_*([C\_N]) = \lambda\kappa \in \langle\kappa\rangle$ is non-zero, so C (the image lies on the κ-line) and D (the image is non-zero) are *both* delivered by the single fact $\lambda \neq 0$.

But the accounting must be honest about *where C's content went*, because the reduction is not as clean as "C becomes a coefficient." The one-dimensionality hypothesis — that the transport-visible persistent sector is one-dimensional and spanned by κ — is *not* an independent input alongside C; it carries most of C's identification content. B (domain rank-one) already forces the image to be cyclic, generated by $\tau_*([C_N])$; that much needs no one-dimensionality. What the one-dimensionality hypothesis adds is precisely that the line the image lands on is *κ's* line rather than some other cyclic direction — which is exactly what C asserts. So Proposition 6.2 does not reduce C to D; it reduces **(C and D) to (the transport-visible sector is one-dimensional and κ-spanned) plus ($\lambda \neq 0$)**, with the κ-identification living in the *first* conjunct, not eliminated. The genuine new content is the compression of two conditions into "one-dimensionality + one non-zero scalar," not the dissolution of C. The guard, marked honestly and consistent with §6.5's accounting below: "find one coefficient and show it is non-zero" is a smaller obligation than "identify an arbitrary class" *only because* the one-dimensionality hypothesis has absorbed the identification — so the reduction relocates C's content into that hypothesis rather than discharging it. With that understood, the residual transport-side question is legible: grant one-dimensionality, and the verdict is the single bit $\lambda \neq 0$.

6.3 C is the minimal identification

Corollary 6.3 (C is the minimal identification) (proven, conditional as in 6.2)

Under A and B, failure of C requires the existence of a second independent transport-visible persistence class, distinct from κ, that $\tau_*([C_N])$ could be. No such class is currently known (Proposition 6.1).

Proof. Under A and B the image is cyclic, generated by $\tau_*([C_N])$. If C fails — $\tau_*([C_N]) \notin \langle\kappa\rangle$ — then the image generator is a transport-visible persistence class not proportional to κ, i.e. a second independent such class distinct from κ. By Proposition 6.1 no such class has appeared in any construction. ■

So C is not an arbitrary identification hypothesis plucked from many candidates; it is the **minimal identification compatible with the observed rank-one persistence structure**. Its only alternative is the existence of a second independent persistence class that the entire programme has, to date, never produced. C is therefore the default, and ¬C is the claim that bears the burden of exhibiting a witness.

6.4 C and D as one scalar question

Proposition 6.2 permits a further compression of the assumption set, and it is worth making explicit, because it changes the *shape* of the transport-side problem.

As originally stated, C and D are two separate transport-side assumptions: C that $\tau_*([C_N]) = \kappa$, and D that $\kappa \neq 0$. But once the transport-visible persistence sector is one-dimensional and spanned by κ, the transported generator has the form $\tau_*([C_N]) = \lambda\kappa$ (Proposition 6.2), and the entire transport-side question becomes a single scalar dichotomy:

$\lambda = 0$ or $\lambda \neq 0$.

The two outcomes are exactly the two faces C and D were separately gesturing at:

- **$\lambda = 0$:** $\tau_*([C\_N]) = 0$, the homological horizon collapses, and the Gate-3 transport branch closes *negatively* — no surviving transport residue.
- **$\lambda \neq 0$:** the transported generator is non-zero and lies on the κ-line; the κ-line is the image of admissible transport, up to normalisation of the generator.

So C and D are not independent once one-dimensionality is granted. They are two presentations of one question:

Does the primitive refinement cycle transport *non-trivially* into the unique persistent κ-direction?

The normalisation caveat, stated precisely. There is one subtlety in equating "$\lambda \neq 0$" with C as literally written. The statement $\tau_*([C\_N]) = \kappa$ is *stronger* than $\tau_*([C\_N]) = \lambda\kappa$. If κ is defined only up to choice of generator of the one-dimensional persistence sector, then any non-zero λ can be absorbed by renormalising κ, and C follows from $\lambda \neq 0$ outright. If instead κ carries a fixed external normalisation (a specific generator, not just a direction), then λ must be checked against that normalisation explicitly, and "$\lambda \neq 0$" gives C's *direction* but not necessarily its literal coefficient. Either way, the essential transport-side verdict is $\lambda \neq 0$: non-zero λ gives the same surviving κ-line, which is what the Gate-3 verdict actually requires; the literal equality $\tau_*([C\_N]) = \kappa$ is then either automatic (free normalisation) or a coefficient-fixing afterthought (fixed normalisation) that does not affect survival.

The compressed assumption set. The transport side therefore reduces:

- *original form:* C and D, two assumptions;
- *compressed form:* $\tau_*([C\_N]) = \lambda\kappa$ with $\lambda \neq 0$, one non-zero-scalar condition;

with the failure condition equally simple — $\lambda = 0$ means no surviving residue, $\lambda \neq 0$ means the κ-line survives. This is the strongest honest version, and the honesty is in the qualifier: the collapse of C and D into one scalar question holds *only after* the rank-one / one-dimensional persistence assumption is granted. Absent one-dimensionality, C and D are genuinely two conditions; with it, they are one. So the compression does not remove an assumption for free — it trades C and D for (one-dimensionality) + ($\lambda \neq 0$), making visible that the transport-side content, once one-dimensionality is assumed, is exactly a single non-zero projection of the domain generator onto the unique persistent κ-direction.

6.5 The picture this assembles

Taken together, §§6.1–6.4 reframe the closure. B is the minimal-rank hypothesis: rank one unless a second generator is exhibited, and none has been. C is the minimal-identification hypothesis: the κ-line unless a second persistence class is exhibited, and none has been. And §6.4 compresses the transport side: granting one-dimensionality, C and D are not two assumptions but one scalar question, $\lambda \neq 0$ — does C_N project non-trivially onto the κ-direction. The conditionality is not eliminated — A's locality, B's (NF), the one-dimensionality guard, and $\lambda \neq 0$ all remain to be settled. But the four assumptions no longer look like four unrelated bets. They look like one rank-one persistence picture, recurring identically across the K_N homology, the σ-mode, the κ-class, and the global closure mode, whose remaining gaps are precisely: confirm no global relation (A), exhibit no second generator / prove (NF) (B), grant one-dimensionality of the transport-visible sector, and pin the single non-zero coefficient λ (C and D jointly). The domain side is A + B; the transport side is one non-zero projection. That is a single emerging structure with heavily-constrained weak points, not a stack of independent hypotheses.

---

7. The Two Disciplines: Two-Sidedness and the Out-of-Scope Hinge

Two disciplines the branch has held throughout must hold at its endpoint too, or the closure would overclaim.

Two-sidedness. The verdict is not a dial running from weak to strong; it has a negative branch, located at D. Discharging A and B bounds the horizon to rank ≤ 1 — which means "a single κ-line *or* nothing." Discharging C identifies the line as κ *if there is a line*. Only D selects between the line and nothing. So "the horizon is ⟨κ⟩" is the *positive* branch of a binary whose *negative* branch (κ = 0, empty horizon, programme dead) is equally consistent with A, B, C. The closure must not be read as "A–D in increasing strength give an increasingly favourable result"; A, B, C give "line or nothing," and D alone gives "line." A reader taking the rank-≤-1 bound as good news has taken half of it.

The out-of-scope hinge. Loop-realisation of γ_D is deliberately *not* among A–D. It is not a condition the closure rests on; it is the question the closure leaves open. A–D close the *homological* horizon — they establish that the transport image is exactly ⟨κ⟩, that $[\gamma_D]$'s class either lies on this line or does not. They do *not* establish that a single reversible merge–split loop *realises* γ_D — that some w has $[\tau_{loops}(w)] = [\gamma_D]$, rather than merely $[\gamma_D]$ lying in the line that loops could in principle reach. That is the loop-level (Stage Two) question, and it is untouched by everything here. The criterion paper's necessary/sufficient asymmetry is the reason: the homological horizon is a necessary condition for realisation and never sufficient. So even full discharge of A–D would settle Stage One completely and Stage Two not at all. Loop-realisation of γ_D remains the final Gate-3 hinge, and it is correctly placed outside the assumption set because it is the thing the conditional closure is honest about *not* delivering.

---

8. What Would Discharge the Closure

Collecting §5 into an ordered programme, the conditional closure becomes unconditional exactly when the following are settled, in roughly increasing difficulty:

1. **Confirm locality (discharges A).** Verify R_MS contains no global period relation. If confirmed, freeness is unconditional and $\mathrm{ord}(C\_N) = \infty$ outright.
2. **Prove (NF) and the bridge (discharges B).** Establish the reversible merge–split normal-form theorem (every closed irreducible a C_N-power), attacking the winding-collection step via an augmented measure; and confirm the presented calculus is Γ_MS. If done, $H_1(\Gamma\_MS) \cong \mathbb{Z}\langle[C\_N]\rangle$ unconditionally.
3. **Establish one-dimensionality O and pin Λ (discharges C and D jointly).** Establish that the transport-visible persistence sector is one-dimensional, spanned by κ (assumption O of §2.1); given O, the transported generator is $\tau*([C\_N]) = \lambda\kappa$ (Proposition 6.2), and the remaining content of both C and D is the single scalar verdict $\lambda \neq 0$ (assumption Λ). Showing Λ ($\lambda \neq 0$) gives the surviving κ-line (D's non-triviality and C's direction together, modulo the §6.4 normalisation caveat); $\lambda = 0$ is the negative branch. So what reads as two assumptions C, D is, given O, one non-zero-projection question Λ — and the canonical transport-side burden is exactly O (a structural claim) plus Λ (one scalar bit).

If all three are discharged, Theorem 3.1 becomes unconditional: the homological horizon is exactly the κ-line, and the Gate-3 verdict of §4 holds outright. And then — the discipline of §7 — one further problem, outside this list, remains: loop-realisation of γ_D, the Stage-Two hinge, which none of 1–3 addresses.

The ordering reflects status: 1 confirms an already-proved-conditional result; 2 proves a reduced-to-concrete property; 3 establishes one-dimensionality and pins the single scalar λ, the genuine open transport-side input — and the sign of λ (zero versus non-zero) is the one that decides the positive-versus-negative branch.

---

9. Conclusion

This paper records the conditional closure of the Gate-3 transport branch. Under the explicit assumption set, as originally framed —

- **A** locality ($\Longrightarrow \mathrm{ord}(C\_N) = \infty$, freeness),
- **B** H_1-transfer ($H_1(\Gamma\_MS) \cong \mathbb{Z}\langle[C\_N]\rangle$, generation),
- **C** $\tau*([C\_N]) = \kappa$ (generator identification),
- **D** $\kappa \neq 0$ (non-triviality) —

the homological horizon of admissible transport is exactly the κ-line,

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

and the Gate-3 verdict follows: native residue, κ, RC_path, and phase-as-memory all survive homologically.

The cleaner statement of what the branch actually rests on (§2.1, §6.4) is the reformulated set **A, B, O, Λ**: domain-side locality (A) and H_1-transfer (B); transport-side one-dimensionality of the persistent sector (O) and non-zero projection $\lambda \neq 0$ (Λ). The transport pair C, D collapses to O + Λ — O carrying the κ-identification, Λ the single non-zero scalar — so the genuine transport-side content is one structural hypothesis plus one bit, not two entangled assumptions. The two presentations are equivalent and the verdict holds under either; A/B/O/Λ is simply the form in which the remaining unknowns are most legible.

The closure is conditional and says so. Its assumptions are annotated with status: A proved under locality, B reduced to the single combinatorial property (NF), C and D open transport-side inputs with D the switch on which the positive verdict turns. And the open pieces are not arbitrary: B is the minimal-rank hypothesis (rank one unless a second generator is exhibited, and none has appeared in any construction), while C and D, given one-dimensionality of the transport-visible sector, reduce jointly to pinning a single non-zero coefficient λ in $\tau_*([C\_N]) = \lambda\kappa$ — so ¬C would require a second independent persistence class the programme has never produced. The four assumptions are therefore one rank-one persistence picture with heavily-constrained weak points, not four unrelated bets. Two disciplines are preserved: the verdict is two-sided (D failing is the negative outcome, not a weaker positive), and loop-realisation of γ_D is out of scope — the Stage-Two hinge the closure deliberately does not deliver.

The purpose is to let the programme move on. The conditionality is explicit enough that no reader could mistake this for a completed proof; equally, recording it prevents the branch from recursing indefinitely on the same transport question. Here is the conditional closure; here are the assumptions; here is the single hinge beyond them. If future work confirms A–D, the homological conclusion follows, and the only question left for Gate-3 is the loop-level realisation of γ_D — at which point this branch hands off to that final hinge, and the programme's next layer can begin.