

The Gate-3 Closure Verdict

From Hub Completion to the Operational Quotient: Existence, Observability, and the Two Substrate Inputs That Remain

Keith Taylor

VERSF Theoretical Physics Programme

General Reader Summary

The transport programme reduced a long question — whether a hidden sevenfold structure leaves any physical residue — first to the topology of the vacuum transport complex Γ_{vac} , then to a single cohomology class $\kappa \in H^1(\Gamma_{\text{vac}}; \mathbb{Z}_7)$. The natural next step seemed to be: build the complex and compute. This paper shows that the verdict splits cleanly into two questions, carries each to its sharpest form, and identifies the one undetermined input that each terminates at.

The first question is *existence*: does the class κ exist at all? This funnels through a single primitive — the rule deciding when a closure hub is complete — because the other admissibility rules are defined only relative to it. The paper shows that a tempting route to existence (reading the substrate's sevenfoldness as a winding that would force the class to live) is foreclosed by the published construction, which is built additively, so the seven lives in the coefficient group and never in a winding. With that closed, existence reduces to one statement about subspaces: the class exists exactly when admissibility *fails* to confine the transport offsets to coboundaries. This replaces "the question is open" with "the question is a single containment," and names the object — the constraint subspace — that decides it.

The second question is *observability*: if the class exists, is it recorded? Here the programme's emergent treatment of time enters. A proposal that a surviving residue is dynamical — that the evolving closure state, not the frozen one, carries it — turns out, when written down, to compute the same class by another name; composing evolution steps around a loop is the definition of the very holonomy the static test already measures. The genuine role of dynamics is the reverse: it is a further filter the class must survive, not a new source for it. Even a class that exists may be coarse-grained into invisibility by the commitment dynamics, because the closure structure is by construction blind to the only observable it carries. Whether that blindness lifts from single steps to whole loops is the observability question, and it is adjudicated by the dynamics, not the static topology.

The two questions terminate at two substrate inputs the foundational constraints deliberately leave open: which offset configurations completion admits, and whether the closure structure's

blindness lifts to loops under commitment. The paper proves the reductions; it does not pretend to the inputs, and is explicit about why they are genuinely unwritten rather than merely unattempted.

Contents

Abstract

1. Introduction and Shared Setup

Part I — Existence: Does the Closure Class Exist?

2. The Dependency Structure of Gate 3
3. The Specification of a Completion Predicate
4. Hub Completion as a Boolean Predicate
5. Conservation and Competition as Derived Rules
6. Sevenfoldness, Shape, and Winding: the Torsion Route Is Foreclosed
7. Occupancy as Coboundary Confinement
8. The Frame-Alignment Reduction
9. The Constraint Subspace as a Selection Input

Part II — Observability: If It Exists, Is It Recorded?

10. Emergent Time and the Moving-Target Proposal
11. The Dynamic Residue Is the Static Class
12. The Memory Escape Is Closed in the Reversible Sector
13. The Operational Quotient
14. Orientation-Blindness and Whether It Lifts to Loops
15. The Form of an Operationally Distinguishable Record
16. The Complete Verdict Chain

Conclusion

Abstract

The transport-group analysis fixed the transport group $G \cong D_7$ with rotation kernel $N = \ker(\pi_{\text{or}}) \cong \mathbb{Z}_7$ and proved single-face rigidity: any surviving closure content is loop-supported. The closure-connection analysis constructed the closure connection, assigned a rotation label $\rho_f \in \mathbb{Z}_7$ to each admissible transport face, detwisted the non-abelian hazard under the Gate-2 coherence condition $H\text{-OC}(\gamma)$, and reduced the survival of a residue to a three-condition chain on $\kappa \in H^1(\Gamma_{\text{vac}}; \mathbb{Z}_7)$: flatness, nontrivial topology, population. The topology paper reduced the topological condition to $\text{rank}_{\mathbb{Z}_7} \partial_2 < \mu$. All reductions terminate at the action of admissibility on the rotation labels.

This paper splits the verdict into existence and observability and sharpens each. **Existence (Part I)**: the input is generated entirely by the hub-completion predicate H , since conservation and competition are defined only on domains H fixes (Proposition 2.1); the published curvature $\Omega(P) = \sum \varepsilon_{\{P,f\}} \rho_f$, additive with unit incidence signs, commits the construction to a *shape* model, so the *winding* mechanism that alone could produce \mathbb{Z}_7 torsion is structurally unavailable and the topology paper's torsion route is unpopulated (Section 6); and occupancy reduces to a containment of subspaces (Proposition 7.1):

Branch A ($\kappa = 0$ on every admissible configuration) $\Leftrightarrow A \subseteq B^1$,

with $A \subseteq C^1(\Gamma_{\text{vac}}; \mathbb{Z}_7)$ the admissible-offset set and $B^1 = \delta C^0$ the coboundaries, graded by the prior flatness containment $A \subseteq Z^1$. The subspace A is generated by completion's frame-alignment rule (Section 8) and is a *selection* input, not a consequence of the foundational constraints (Section 9). **Observability (Part II)**: the proposal that a surviving residue is dynamical collapses, because composing per-step evolution around a loop is the definition of holonomy and $\Delta\rho = (U_\gamma - I)\rho$ equals the static class (Proposition 11.1); the only escape, non-Markovian memory, cannot reach the reversible sector where the holonomy lives (Proposition 12.1). The genuine dynamical content is the operational quotient $N_{\text{hol}} \rightarrow N_{\text{obs}}$ (Definition 13.1), a reduction and not a source (Proposition 13.2): a class is observable only if commitment dynamics record it, which — since the kernel is orientation-blind by construction — requires per-face orientation-blindness to fail to lift to loops (Proposition 14.1); the form such a record would take couples the reversible kernel holonomy to an irreducible committed quantity, and that coupling is exactly $N_{\text{obs}} \neq 0$ (Section 15).

Markers: the factorisation, the shape/winding distinction, the coboundary reduction, the equality $\Delta\rho = \kappa$, and the reversible-sector argument are [proven]; the operational-quotient framing is [proven as a reduction]; the predicate H and subspace A , and whether orientation-blindness lifts, are [conjectural / open]; the lean toward Branch A and toward $N_{\text{obs}} = 0$ from programme history is [heuristic].

1. Introduction and Shared Setup

The transport programme produced a chain of objects,

$$G \rightarrow N \rightarrow \rho \rightarrow \kappa,$$

the transport group, its orientation-blind kernel, the rotation labels on transport faces, and the candidate closure residue. The closure-connection paper reduced the survival of κ to a three-condition chain on $H^1(\Gamma_{\text{vac}}; \mathbb{Z}_7)$ — flatness, then nontrivial topology, then population — and the topology paper reduced the topological condition to $\text{rank}_{\mathbb{Z}_7} \partial_2 < \mu$. The whole apparatus terminates at one undetermined input: how admissibility acts on the rotation labels.

This paper organises the remaining work into two questions and carries each to its sharpest statement.

Existence. Does the class κ exist — does admissibility leave a nonzero class in $H^1(\Gamma_{\text{vac}}; \mathbb{Z}_7)$? (Part I.)

Observability. If κ exists, is it operationally recorded — does it survive the coarse-graining the commitment dynamics impose? (Part II.)

The two are independent terminators: existence governs the static chain (hub completion, coboundary confinement); observability governs the operational quotient $N_{\text{hol}} \rightarrow N_{\text{obs}}$ left open by the transport-group paper. A residue survives only if both clear.

1.1 Shared setup (used throughout)

The following are imported from the published corpus and used in both parts.

- **The closure connection.** Each admissible oriented face f carries a transport element $g_f \in D_7$ with rotation coordinate $\rho_f \in \mathbb{Z}_7$. Under $H\text{-OC}(\gamma)$ the connection is honest- \mathbb{Z}_7 -valued and $\rho = \{\rho_f\}$ is a 1-cochain $\rho \in C^1(\Gamma_{\text{vac}}; \mathbb{Z}_7)$. The curvature is the additive cochain

$$\Omega(P) = \sum_{\{f \subset \partial P\}} \varepsilon_{\{P, f\}} \rho_f \pmod{7}, \quad \varepsilon_{\{P, f\}} \in \{+1, -1\},$$

with ε the honest cellular incidence signs. Write $d^0 : C^0 \rightarrow C^1$, $d^1 : C^1 \rightarrow C^2$; $B^1 = \text{im } d^0$ (coboundaries), $Z^1 = \ker d^1$ (cocycles), $H^1 = Z^1/B^1$, and $B^1 \subseteq Z^1$ since $d^1 d^0 = 0$. By Theorem 5.3 of the closure-connection paper a flat connection defines $\kappa = [\rho] \in H^1$, with $[\rho] = 0 \Leftrightarrow \rho \in B^1$. The gauge shifts are exactly $\text{im } d^0$: a coboundary is a pure relabelling.

- **Single-face rigidity.** A single face carries no orientation-invariant closure content (transport-group paper, Theorem 3.1); the kernel is gauge on one face, so any residue is loop-supported.

- **The reduction chain.** $N_alg \cong \mathbb{Z}_7$ (the algebraic kernel) $\rightarrow N_hol$ (the part surviving to loop holonomy, i.e. κ) $\rightarrow N_obs$ (the part surviving operational observation). Part I concerns $N_alg \rightarrow N_hol$; Part II concerns $N_hol \rightarrow N_obs$.
- **The reversible/dissipative split.** The kernel acts by orientation-blind rotation $x \mapsto x + c$, which preserves transition probabilities; by Axiom A3 of the Hamiltonian-as-generator paper such evolution is reversible and unitary. Irreversible evolution decomposes into a reversible Hamiltonian component and a dissipative Lindblad component, with the Markovian restriction assumed and non-Markovian effects confined to the dissipator.

Part I — Existence: Does the Closure Class Exist?

2. The Dependency Structure of Gate 3

The programme's state is the chain

$$H \rightarrow \Gamma_vac \rightarrow \partial_2 \rightarrow H^1(\Gamma_vac; \mathbb{Z}_7) \rightarrow \kappa,$$

with H the hub-completion predicate. The published papers supplied the three rightmost arrows; the leftmost — from H to Γ_vac — is the one no paper has supplied, and every object downstream is symbolic until it is.

Proposition 2.1 (Factorisation through completion). [proven, from the definitions of the three rules]

The admissibility map $\mathcal{A} : \mathcal{F} \rightarrow \{\text{cellular 2-complexes}\}$ factors as a strict composition

$$\mathcal{A} = \mathcal{C}_comp \circ \mathcal{C}_cons \circ \mathcal{C}_complete,$$

each stage taking the output of the previous as its domain. Hence \mathcal{A} is determined by the completion predicate H together with two rules whose domains are fixed by H .

Argument. Completion $\mathcal{C}_complete = H$ is a predicate on individual frames; its 1-set is the admissible hubs, the 0-cells. Conservation \mathcal{C}_cons is a predicate on candidate faces *between completed hubs* — its domain is pairs H admits, so it cannot be posed on a pair H has not admitted. Competition \mathcal{C}_comp resolves incompatibilities among *conserved* relations — its domain is the faces \mathcal{C}_cons passes. Each stage is undefined until its predecessor acts, which is strict composition. Conservation and competition contribute no data independent of H : they prune within a candidate set H alone generates.

So H is the unique upstream bottleneck: no computation involving Γ_vac , ∂_2 , or H^1 can begin while H is unspecified.

3. The Specification of a Completion Predicate

"Completed closure" carries no definition. A definition must answer four questions, here converted into conditions on H.

1. **What is a boundary relation?** — the type of object a frame's relations are.
2. **What operation performs closure?** — the operation closing six relations through a seventh.
3. **What distinguishes the seventh relation?** — the asymmetry making completion sevenfold.
4. **What constitutes an unresolved distinction?** — the obstruction whose absence certifies completion.

Questions 1–2 are *constructive* and require substrate primitives not supplied here (Section 9). Questions 3–4 are *structural* and constrain the form of H independently of 1–2.

4. Hub Completion as a Boolean Predicate

Definition 4.1 (Completion predicate)

Let \mathcal{F} be the set of candidate $K = 7$ closure frames. The **completion predicate** is

$H : \mathcal{F} \rightarrow \{0, 1\}$, $H(F) = 1 \Leftrightarrow F$ is a completed closure hub,

with $H^{-1}(1)$ the 0-cells of Γ_{vac} . [conjectural — specified, not constructed]

Definition 4.2 (Well-formedness conditions)

H is an **admissible completion predicate** only if it satisfies: **(W1) Locality** — $H(F)$ depends only on F and its immediate closure data, making Γ_{vac} constructible bottom-up; **(W2) Decidability** — H is total and computable on finite data; **(W3) Sevenfold reference** — H certifies completion through six relations closed by a seventh, per the $K = 7$ architecture; **(W4) Distinction-freeness** — $H(F) = 1$ requires the absence of an unresolved distinction, completion being the certified vanishing of that obstruction rather than the mere presence of seven relations.

By Proposition 2.1, once H is specified the rest of admissibility is derived; constructing H is the correct next objective.

5. Conservation and Competition as Derived Rules

Conservation acts on candidate faces between completed hubs — domain fixed by H — deciding which carry an admissible transport face (orientation parity after Gate 2, closure-frame compatibility, non-duplication of distinguishability content); output: the 1-cells. **Competition** acts on collections of conserved relations — domain fixed by conservation — selecting a maximal compatible subset and determining which minimal cycles survive as plaquettes; output: the 2-cells. The order completion \rightarrow conservation \rightarrow competition is forced by where each rule is defined, not chosen. [both conjectural — domains fixed, contents to be specified]

6. Sevenfoldness, Shape, and Winding: the Torsion Route Is Foreclosed

6.1 The neutrality requirement

H inherits the circularity guard: it must be definable from the closure rules alone and *capable* of returning either $H^1 = 0$ or $H^1 \neq 0$, never gerrymandered to predetermine survival. One might fear (W3) conflicts with this: the substrate is sevenfold, and the topology paper's Proposition 4.2 showed integral 7-torsion lands in the \mathbb{Z}_7 sector even when rational homology vanishes. If sevenfold completion produced torsion as a matter of course, (W3) would force survival. The published construction dissolves the fear by having already chosen the mechanism.

6.2 Shape is not winding

"Sevenfold completion" is ambiguous between two attaching mechanisms, only one producing torsion.

Proposition 6.1 (Polygonal length 7 does not force survival). [proven, cellular homology] *The filled heptagon — seven hubs, seven faces $e_j = (v_j \rightarrow v_{j+1} \text{ mod } 7)$, one plaquette $\partial_2 P = e_0 + \dots + e_6$ — has $|V| = |E| = 7$, $|F| = 1$, $\mu = 1$, $\text{rank } \partial_2 = 1$, so $\dim H_1 = 0$. A sevenfold-shaped cell gives $H_1 = 0$.*

Proposition 6.2 (Winding degree 7 forces survival). [proven, cellular homology] *One hub, one loop edge a , one plaquette $\partial_2 P = 7a$: $\mu = 1$, over \mathbb{Z} the image is $7\mathbb{Z}$ so $H_1(\mathbb{Z}) = \mathbb{Z}/7$, and over \mathbb{Z}_7 the entry reduces to 0, $\text{rank } \partial_2 = 0$, $\dim H_1 = 1$.*

A primitive period-7 orbit through distinct states realises the heptagon (shape, torsion-free); only a degree-7 attaching map carries $\mathbb{Z}/7$.

6.3 The published construction commits to shape

The closure-connection curvature $\Omega(P) = \sum \varepsilon_{\{P,f\}} \rho_f$ has each face entering *additively* with a *unit* coefficient. Winding — a face entering with coefficient 7, an attaching degree — has no slot in this cochain. Equivalently, $\kappa = d^1\rho = 0 \pmod{\text{im } d^0}$ over the field \mathbb{Z}_7 is $\text{Hom}(H_1, \mathbb{Z}_7)$, which sees only free rank and never torsion.

Consequence. The torsion availability of the topology paper's Proposition 4.2 is valid mathematics but *structurally unpopulated*: the kernel acts by free translation (the shape pole) and the incidence signs are units. Whatever H is, it feeds the additive cochain and cannot reintroduce winding without rewriting Ω . Neutrality is therefore preserved not by a shape/winding fork — that fork is closed — but by a different and genuinely open datum: whether admissibility confines the offsets to coboundaries. That is the subject of Section 7.

7. Occupancy as Coboundary Confinement

This is the central result of Part I.

Proposition 7.1 (Occupancy = coboundary confinement). [proven, given the closure-connection construction and standard cellular cohomology]

Let $A \subseteq C^1(\Gamma_{\text{vac}}; \mathbb{Z}_7)$ be the admissible-offset set imposed by completion, conservation, and competition. Under $H\text{-OC}(\gamma)$ the admissible configurations fall into exactly three cases, refining the closure-connection paper's Proposition 6.1:

1. **$A \subseteq B^1$.** Every admissible offset is a coboundary, hence flat and gauge-trivial, so $\kappa = 0$ for every admissible configuration. **Branch A** in its strongest form — the closure channel is pure redundancy.
2. **$A \subseteq Z^1$ but $A \not\subseteq B^1$.** Admissibility forces flatness but not confinement: every admissible ρ is a cocycle, some admissible cocycle is not a coboundary, so κ can be nonzero, with surviving classes $A/(A \cap B^1) \subseteq H^1$. **Branch B realised** for those configurations.
3. **$A \not\subseteq Z^1$.** Admissibility does not force flatness: some admissible ρ has $\Omega(P) \neq 0$ on a plaquette, so κ is undefined as a class; the obstruction lives in the curvature (failure of condition (1) of Proposition 6.1).

Argument. A flat connection defines $\kappa = [\rho] \in Z^1/B^1$ with $[\rho] = 0 \Leftrightarrow \rho \in B^1$. Case 1: $A \subseteq B^1 \subseteq Z^1$ gives flatness and zero class. Case 2: $A \subseteq Z^1$ gives flatness, $A \not\subseteq B^1$ gives a nonzero class, with admissible image $A/(A \cap B^1)$. Case 3: $A \not\subseteq Z^1$ gives $d^1\rho \neq 0$, representing no class. The cases are exhaustive and exclusive by whether $A \subseteq B^1$, $A \subseteq Z^1 \not\supseteq B^1$, or $A \not\subseteq Z^1$.

The verdict is now two nested containments — is $A \subseteq Z^1$ (flatness), and if so is $A \subseteq B^1$ (non-occupancy) — each a finite rank computation once A is known.

Why the generic count is uninformative. Of $7^3 = 343$ offset triples on a triangle, 49 sum to zero and 294 do not, suggesting "6 of 7 are irreducible." This is a fact about the *unconstrained* space C^1 and decides nothing: the same 294/343 holds for a pure gauge field $\rho = \delta x$, which is

entirely Branch A — the 294 non-zero-sum triples are simply not admissible. Genericity in C^1 is not genericity in A. Only the position of A decides the verdict, which is why Proposition 7.1 is a containment of A and not a count.

8. The Frame-Alignment Reduction

A is generated, face by face, by completion's **frame-alignment rule**: given two completed hubs sharing a face, which slot of one maps to which slot of the other determines g_f , hence ρ_f , hence A. By single-face rigidity the verdict is decided not on single faces but on loops, with minimal support a 3-hub triangle. The decisive property is:

Is the inter-frame offset ρ_f **decomposable** — $\rho_f = a(\beta, \cdot) - a(\alpha, \cdot)$ for some 0-cochain a — or **indecomposable**, a primitive relative rotation with no per-side splitting?

Decomposable $\Rightarrow \rho = \delta a \in B^1 \Rightarrow A \subseteq B^1 \Rightarrow$ Branch A. Indecomposable $\Rightarrow \rho$ need not be a coboundary \Rightarrow Branch B possible. The triangle test reads it off: the composite of the three alignments around $\alpha \rightarrow \beta \rightarrow \gamma \rightarrow \alpha$ is the identity (decomposable, Branch A) or a nontrivial rotation s^k (irreducible curvature, Branch B foothold).

The trap, stated once. A slot depending on the *face*, or on the *pair* (h, h') , is still a well-defined cochain and its difference is still a coboundary — still Branch A. A relational ontology is *necessary* but not *sufficient*: relational-but-consistent comparisons descend from a common reference and decompose; only relational-*and-frustrated* comparisons — admitting no global slot-numbering, no common reference — give a non-coboundary. The substrate fork: does a common reference frame exist (every offset a difference against it — Branch A), or do frames compare only pairwise with no common standard, so the comparisons can fail to close around a loop (Branch B)? Indecomposability is strictly stronger than relationality and must be checked against the triangle composite, not its verbal plausibility.

9. The Constraint Subspace as a Selection Input

Part I reduces existence to the position of A and gives A a concrete generator. It does not compute A, and there is a structural reason this is not routine deferral.

The open content is questions 1–2 of Section 3: the nature of a boundary relation [open, requires substrate input], and the closing operation with its alignment rule [open, requires substrate input] — the latter being precisely whether the offset is decomposable (Branch A) or indecomposable (Branch B). A methodological observation sharpens the status. The foundational admissibility framework (finite distinguishability FD, irreversible commitment IC) is *constraint-complete but not dynamically complete*: FD and IC fix the *form* any law must take but do not *select* among admissible laws. Coboundary confinement — which admissible offset configurations occur — is

a *selection* question, and both branches respect FD and IC. So A is not deducible from the foundational constraints; it is a selection input the framework explicitly leaves to be supplied. This is why every attempt to derive the verdict from general features (sevenfoldness, relationality, topology, evolutionary dynamics) has been neutral or has collapsed to Branch A: those are form-level features, and the verdict is selection-level.

A [heuristic] lean: programme history — the Born rule, the transport group, the gauge reduction, single-face rigidity — repeatedly shows freedom collapsing to constraint to uniqueness, which would favour $A \subseteq B^1$, Branch A. This is the strongest evidence available, but it is *form-level* collapse applied to a *selection-level* question, where the same framework declines to collapse (leaving gauge groups, matter content, couplings open). The lean is real but weaker than the bare pattern suggests; the fork is recorded as open.

Part II — Observability: If It Exists, Is It Recorded?

Part I governs whether κ exists. Part II governs whether an existing κ is operationally recorded — the arrow $N_{\text{hol}} \rightarrow N_{\text{obs}}$. Because VERSF treats time as emergent from commitment, one might expect the dynamics to bear on existence; Sections 10–12 show they do not, and Sections 13–15 show where they genuinely bear.

10. Emergent Time and the Moving-Target Proposal

The proposal observes accurately that every static construction has collapsed into a coboundary, and infers that the residue may be *dynamical*: that the static loop $\alpha \rightarrow \beta \rightarrow \gamma \rightarrow \alpha$ should be read as an evolving loop $\alpha(t_0) \rightarrow \beta(t_1) \rightarrow \gamma(t_2) \rightarrow \alpha(t_3)$ whose endpoint occupies the same network position but need not be the same closure state. It posits a TPB evolution operator U , a loop composite $U_\gamma = U_n \cdots U_1$, and a dynamic residue $\Delta\rho = (U_\gamma - I)\rho$, with the loop dynamically closed when $U_\gamma\rho = \rho$; survival is claimed to live in $\Delta\rho$ where the static offsets vanish.

Two remarks. The inference "static collapses, therefore dynamical residue" converts every prior refutation into support and so must be tested on the mathematics, not accepted on the pattern it cites. And the proposal is right that emergent time should bear on Gate 3 — the question is *where*.

11. The Dynamic Residue Is the Static Class

Proposition 11.1 (Holonomy equals loop composition). [proven, definition of holonomy and the closure-connection construction]

If the per-step evolution acts on the labels by the closure transport, the step across f carrying rotation by ρ_f , then the loop composite acts by the signed loop sum,

$$U_\gamma : \rho \mapsto \rho + \sum_{\{f \in \gamma\}} \varepsilon_{\{\gamma, f\}} \rho_f,$$

so $\Delta\rho = (U_\gamma - I)\rho$ vanishes around every loop $\Leftrightarrow \sum \varepsilon \rho \equiv 0 \pmod{7}$ for every $\gamma \Leftrightarrow \kappa = 0$.

Argument. Transport composes by multiplying group elements, hence adding rotation coordinates with incidence signs; the loop composite is translation by $\sum \varepsilon \rho$, which is Ω on γ . By Theorem 5.3 of the closure-connection paper this is the holonomy of the flat connection, vanishing on all loops iff $\kappa = 0$.

The content is structural: composing local steps around a loop *is* the definition of holonomy, and \mathbb{Z}_7 holonomy is classified by $\text{Hom}(\pi_1, \mathbb{Z}_7) = H^1$. "Dynamic residue" is κ renamed; "dynamic closure" is flatness. The static and dynamic tests are one test, because U_γ on the labels is exponentiated $\sum \varepsilon \rho$. This explains why the collapse recurs: any object built from local steps composed around loops lands in H^1 .

12. The Memory Escape Is Closed in the Reversible Sector

Proposition 11.1 leaves one escape: U_γ must *not* factor into local steps, requiring genuine memory in the evolution.

Proposition 12.1 (Memory cannot reach the reversible holonomy). [proven, given the GKSL decomposition]

The closure transport is reversible structure; non-Markovian memory is confined to the dissipative sector and does not alter the reversible holonomy. Hence U_γ on the labels is a product of memoryless local factors regardless of any memory in the dynamics, and $\Delta\rho = \kappa$ stands.

Argument. The kernel acts by orientation-blind rotation, preserving transition probabilities, hence reversible and unitary (Axiom A3). Irreversible evolution decomposes into reversible Hamiltonian and dissipative Lindblad components; the Markovian restriction is assumed, and non-Markovian effects modify only the dissipator while leaving the reversible component structurally untouched. Memory kernels are dissipative corrections; the closure holonomy does not occupy that sector, so the reversible loop composite remains a product of local factors and its holonomy is flat-connection holonomy.

The walking-round-the-Earth analogy reinforces this: the "frame has moved" is *real displacement* — observable, committed motion — which is the *dissipative* sector, not the reversible holonomy. If the loop commits genuine record the residue is not $\kappa \in H^1$; if it commits none the traveller returns up to gauge and $\Delta\rho$ is a coboundary. The analogy points away from a dynamic closure residue.

13. The Operational Quotient

Dynamics supplies no new source of residue, but it supplies a real and previously open *reduction*.

Definition 13.1 (Operational quotient)

A holonomy class $\kappa \in N_{\text{hol}}$ is *operationally observable* if some admissible irreversible-commitment process produces a committed record differing between a loop carrying κ and one carrying the trivial class. Classes are *operationally equivalent*, $\kappa \sim_{\text{op}} \kappa'$, if no admissible committed record distinguishes them. The observable sector is $N_{\text{obs}} = N_{\text{hol}} / \sim_{\text{op}}$.

Finite distinguishability makes this quotient nontrivial in principle: committed records have bounded resolution and the distinguishable outcomes are finite, so $N_{\text{hol}} \rightarrow N_{\text{obs}}$ may have nonzero kernel. A class in that kernel is *real but unobservable* — present in the reversible holonomy, absent from every committed record. The committing process is the tick dynamics — irreversible commitment, the source of emergent time — so whether a class produces a distinguishable record is a question about the dynamics, not the static cohomology.

Proposition 13.2 (Dynamics is a reduction, not a source). [proven as a reduction]

N_{obs} is a quotient of N_{hol} , so $N_{\text{obs}} \neq 0$ requires $N_{\text{hol}} \neq 0$: the dynamics can only *reduce* the surviving sector, never enlarge it. Even granting Branch B ($N_{\text{hol}} \neq 0$), N_{obs} may be zero; and $N_{\text{hol}} = 0$ forces $N_{\text{obs}} = 0$ a fortiori. The corrected role of emergent time is the inverse of the proposal: not a dynamical source for a residue the static test misses, but a dynamical sieve an existing residue must pass.

14. Orientation-Blindness and Whether It Lifts to Loops

The operational quotient has a sharp form, because the kernel is defined as an invisibility. $N = \ker(\pi_{\text{or}})$ is the orientation-blind part of transport — the elements trivial on orientation, the only observable transport carries at the face level. Single-face rigidity states a single face carries no orientation-invariant content. The hope of any residue is that *loops* resurrect what faces hide: κ a loop-level observable though every face is invisible.

Proposition 14.1 (Observability requires orientation-blindness to fail at loop level). [open, dynamical — the precise content of $N_{\text{hol}} \rightarrow N_{\text{obs}}$]

$N_{\text{obs}} \neq 0$ requires an admissible loop γ on which the commitment dynamics produce an orientation-distinguishable committed record, even though the kernel is orientation-blind on every face of γ . Equivalently: $N_{\text{obs}} \neq 0$ iff per-face orientation-blindness fails to lift to per-loop orientation-blindness under the committing dynamics.

If blindness lifts, every class commits the same record as the trivial class, \sim_{op} identifies all of N_{hol} , and $N_{\text{obs}} = 0$ — the kernel stays invisible at every scale, the channel operationally

redundant whether or not topologically occupied. If blindness does not lift, $N_{\text{obs}} \neq 0$ and the residue is observable. The adjudicator is the tick dynamics, not the static class; this is the one substantive way dynamics enter, and it is genuinely open.

A caution consistent with the programme: the burden leans toward the lift *holding* ($N_{\text{obs}} = 0$), since orientation-blindness is the defining property of the kernel and is exhibited per-face by Theorem 3.1; for it to fail at loop level the dynamics would have to manufacture an orientation observable from a composition of orientation-invisible steps — the operational analogue of producing a non-coboundary from coboundary-like local data. That is the demanding side of the fork, and the heuristic evidence is the programme's history of such resurrections failing.

15. The Form of an Operationally Distinguishable Record

Definition 13.1 fixes N_{obs} mathematically but leaves it abstract. This section states what a distinguishing record would have to *look like* — as a specification of the target, not a demonstration that the target is met. By construction any such record is an instance of $N_{\text{obs}} \neq 0$, the demanding side of Proposition 14.1; the value here is to expose what it would require, and the requirement is the open question.

Two example-forms exhaust the natural candidates, and each comes with a wall.

The interferometric form. Two admissible loops with identical orientation outcomes face-by-face, differing only in a relative offset that no sub-path reveals but that commits a distinguishable record when the loop closes — the way a closed interferometer registers an enclosed flux invisible on any segment. This is the honest shape of " κ as a loop-only observable." *The wall:* the enclosed quantity is valued in the kernel, which is orientation-blind, so the committed record would have to register a kernel-valued accumulation that the orientation observable cannot see per-face yet can see per-loop. Whether the commitment dynamics close a kernel-valued accumulation into an orientation-distinguishable record is exactly Proposition 14.1, restated. The form makes the requirement vivid; it does not satisfy it.

The commitment-counting form. Two admissible loops with the same orientation at every face but a different number or pattern of irreversible commitments accumulated around the traversal, so the committed ledger differs though the reversible observations agree. This is the more VERSF-native candidate. *The wall:* a difference in committed count is a difference in the *dissipative* sector, and Proposition 12.1 placed κ in the *reversible* sector. So if the loops genuinely differ in commitment, the distinguishing quantity is not κ but an irreducible committed record — and observability of κ then requires that record to be *correlated with* κ rather than independent of it. If independent, it witnesses nothing about the kernel holonomy. If correlated — if the commitment pattern around a loop is forced to track the kernel holonomy — that correlation is what $N_{\text{obs}} \neq 0$ demands, and establishing it is the open work.

Both forms converge on the same structural requirement: a distinguishing record must *couple* an orientation-blind, reversible quantity (κ) to a committed, observable one. That coupling is precisely the failure of orientation-blindness to lift — the reversible kernel holonomy leaving an irreducible mark in the irreversible ledger. Its existence is $N_{\text{obs}} \neq 0$; its absence is $N_{\text{obs}} = 0$.

Nothing in this section establishes the coupling. It states the form the coupling must take, which is the most that can be said before the commitment dynamics around a loop are specified — the observability analogue of the constraint subspace A being unwritten in Part I.

16. The Complete Verdict Chain

Both parts assembled, the full chain to an observable closure residue is

$H\text{-OC}(\gamma)$ on $\Gamma_{\text{vac}} \rightarrow \text{flatness } (A \subseteq Z^1) \rightarrow \text{occupancy } (A \not\subseteq B^1, \text{ so } N_{\text{hol}} \neq 0) \rightarrow \text{observability}$
(orientation-blindness fails to lift, so $N_{\text{obs}} \neq 0$),

four conditions in order, each a potential terminator. Part I (hub completion, coboundary confinement) governs the first three; Part II (operational quotient) governs the fourth. The moving-target proposal does not add a fifth channel — it *is* the fourth condition, mislocated as a source; there is no separate dynamic residue, since $\Delta\rho = \kappa$.

The structure is uniformly subtractive: every condition is a containment or vanishing the residue must avoid, and the dynamics add one more hurdle rather than a bypass. The honest reading of the recurring collapses is not that the residue hides in the dynamics. It is that holonomy-equals-loop-composition forces local constructions into H^1 (Proposition 11.1); that the closure connection's additivity forecloses the torsion route (Section 6); that coboundary confinement is a selection question the foundations leave open (Section 9); and that, even past these, the operational quotient stands as a final sieve adjudicated by the tick dynamics (Sections 13–15).

The verdict terminates at two substrate inputs, neither in the published corpus and neither deducible from the foundational constraints:

- **(Existence)** the constraint subspace A — whether completion's frame-alignment rule makes the offsets decomposable (Branch A) or indecomposable (Branch B);
- **(Observability)** whether orientation-blindness lifts to loops under the commitment dynamics — whether the reversible kernel holonomy couples to the irreversible ledger.

Conclusion

The Gate-3 closure verdict splits into existence and observability, and this paper reduces each to a single open input while proving everything between.

Existence reduces to coboundary confinement: with the winding route foreclosed by the published additive construction (Section 6), the class exists exactly when admissibility fails to confine the transport offsets to coboundaries, $A \not\subseteq B^1$ (Proposition 7.1), a verdict generated by completion's frame-alignment rule and turning on whether the inter-frame offset is decomposable

against a common reference (Section 8). Observability reduces to the operational quotient: the proposal of a dynamical residue collapses because composing evolution around a loop is the definition of the static holonomy, $\Delta\rho = \kappa$ (Proposition 11.1), with the memory escape closed in the reversible sector (Proposition 12.1); the genuine dynamical content is whether an existing class survives committing into an observable record, which requires the kernel's defining orientation-blindness to fail to lift from faces to loops (Proposition 14.1), a coupling between the reversible holonomy and the irreversible ledger whose form is specified but whose existence is open (Section 15).

The reductions are proven. The two remaining inputs — the constraint subspace A , and the lifting of orientation-blindness under commitment — are *selection* and *dynamical* facts respectively, both left open by the foundational constraints, which fix form and not content. Programme history leans toward Branch A and toward $N_{\text{obs}} = 0$, but the lean is form-level and heuristic, applied to selection-level and dynamical questions, and is recorded as suggestive rather than decisive. The verdict is not delivered; it is reduced to exactly the two facts the substrate must supply, and the path from those facts to the verdict is now mechanical.