

The Transport Group of the $K = 7$ Closure

Constructing G from the Closure-Frame Primitives — D_7 , the \mathbb{Z}_7 Channel, and Why the Surplus Lives on Loops and Not on Faces

Keith Taylor VERSF Theoretical Physics Programme versf-eos.com

General Reader Summary

A companion paper reduced a central question — does the structure above Gate-2 carry one obstruction or two? — to a single fact about how transport across a face is built: does transport carry the substrate's sevenfold closure structure, or only an orientation sign? That paper proved everything around the fact but could not answer it, because answering it needs the definition of transport itself. This paper supplies that definition and computes the answer as far as it can be computed.

The picture is this. A hub is a small closure cell: six boundary channels that return through a seventh, completing channel, arranged in a cyclic order — a sevenfold *closure frame*. A shared face is where two neighbouring hubs compare their frames, and transport across the face is the rule that re-expresses one hub's frame in the other's. Orientation is just the read-direction of that re-expression: same sense, or reversed.

Once transport is "re-express the sevenfold frame," its symmetry group is forced. The symmetries of a sevenfold cyclic order are exactly the rotations (which keep the cyclic sense) and the reflections (which reverse it) — the dihedral group of the heptagon, with fourteen elements. Orientation is the rotation-or-reflection distinction, and the part orientation throws away — the closure channel — is exactly the seven rotations. So if transport carries the frame at all, the closure channel is a clean sevenfold structure, and it cannot be disguised as orientation, because reversing the read-direction is the very same act as reversing the cycle. The weaker alternative — transport forgetting the frame and recording only the sign — gives a transport group with two elements and an empty channel; we keep it on the table but note it makes the defining seven of the construction idle, which sits badly with calling it a sevenfold closure.

Then comes the surprise the computation forces, and it is worth stating plainly because it corrects what a companion paper had guessed about *where* the sevenfold freedom would show up. One might expect it to appear already when two hubs meet across a single face. It does not. A single face has too much relabelling freedom at its two ends: whatever sevenfold twist transport puts there can always be absorbed by re-choosing the local frames, leaving only the orientation sign behind. A single face, in other words, is *pure convention* in everything except its orientation. The sevenfold content is real, but it cannot be seen on one face; it can only be seen by going around a

closed loop and asking whether the frame comes back to itself or returns rotated. That is a loop quantity, not a face quantity.

This relocates the whole surplus. The companion paper had guessed that if the channel were populated, the surplus would appear already at the pairwise (single-face) level and would undercut a third paper's claim that everything pairwise is settled by Gate-2. The construction shows the reverse: pairwise *is* settled by orientation alone, that third paper's claim is vindicated, and the sevenfold surplus — if it is there — lives one level up, as a closure holonomy around loops, measured in sevenths of a turn and coupled to orientation. What remains genuinely open is whether the substrate's own admissibility rules ever produce a nonzero such holonomy, or force it to vanish. That is now a sharp, well-posed question with the same shape as Gate-2 itself, one notch higher — and it is the honest next step.

The effect, taken together, is that the question itself has changed shape. It began as "what is the transport group?" — and that is now answered: the group is fixed, and its hidden, orientation-blind part is a clean sevenfold channel. The live question is no longer what the channel *is* but whether any of it *survives being observed* — whether, after the substrate's own rules and the freedom to relabel frames have had their say, anything in that channel leaves a trace that a measurement could see. The seven is there in the structure; whether it is there in the physics is what remains.

Abstract

A companion paper localized the obstruction-count above Gate-2 to one fact: whether face-transport carries the \mathbb{Z}_7 closure or only orientation. This paper constructs the transport group G from the closure-frame primitives, settles its isomorphism type under each modelling reading, and derives the structural location of the surplus κ — correcting the companion paper's branch-(N) reading in the process.

The primitives. A hub is a local $K = 7$ closure frame: six boundary channels closing through a seventh, carrying a cyclic order with labels in \mathbb{Z}_7 and an orientation (read-direction) sign. A shared face compares the frames of two adjacent hubs; transport across it re-expresses one hub's frame in the other's; orientation is the parity of that re-expression. The carried object is the sevenfold cyclic closure frame.

The transport group (Theorem 2.2, [Proven under the strong model]). The structure group of a sevenfold *cyclic-successor* frame is the dihedral group

$$G \cong D_7 = \langle s, r \mid s^7 = r^2 = e, r s r^{-1} = s^{-1} \rangle, |G| = 14,$$

with rotations $\langle s \rangle \cong \mathbb{Z}_7$ preserving the cyclic sense and reflections $r s^k$ reversing it. The larger affine group $AGL(1, 7) = \mathbb{Z}_7 \rtimes \mathbb{Z}_6$ is excluded: dilations $x \mapsto ax$ with $a \neq \pm 1$ do not preserve the

closure cycle's successor relation (Remark 2.4). The orientation projection is the rotation/reflection parity

$$\pi_{\text{or}}: D_7 \rightarrow \mathbb{Z}_2, s^k \mapsto 0, r s^k \mapsto 1, \ker(\pi_{\text{or}}) = \langle s \rangle \cong \mathbb{Z}_7.$$

So the closure channel is exactly the rotation subgroup \mathbb{Z}_7 .

D_7 is forced over \mathbb{Z}_{14} (Theorem 2.5, [Proven under the strong model]). The companion paper's residual \mathbb{Z}_{14} -vs- D_7 ambiguity is resolved by the primitives. Orientation is the read-direction of the frame, and the read-direction of a cyclic closure *is* its cyclic sense; reversing orientation therefore reverses the cyclic order, which is inversion $s \mapsto s^{-1}$. Hence the orientation \mathbb{Z}_2 acts on \mathbb{Z}_7 by inversion, giving D_7 . The abelian alternative $\mathbb{Z}_{14} = \mathbb{Z}_7 \times \mathbb{Z}_2$ would require orientation-reversal to fix the cyclic labels — to reverse read-direction without reversing the cycle — contradicting the identification of orientation with cyclic sense. \mathbb{Z}_{14} is thus available only under a reading that decouples orientation from the frame, which the strong model does not permit.

The modelling fork (Proposition 2.6, [Proven]). Under the *weak* reading (transport records only induced boundary orientation, forgetting closure labels), $G \cong \mathbb{Z}_2$, $\ker(\pi_{\text{or}}) = 1$, and rigidity holds outright — but the construction's defining seven is then idle in transport. Under the *strong* reading (transport carries the closure frame, the natural reading of a closure construction), $G \cong D_7$, $\ker(\pi_{\text{or}}) = \mathbb{Z}_7$, and $7 \mid |G|$, so the closure channel is structurally nontrivial. The companion paper's Proposition 6.3 then applies: the seven, once transported, is forced into the kernel.

Single faces are pure gauge — the surplus is a loop quantity (Theorem 3.1, [Proven under the strong model]). This is the principal structural result, and it corrects the *level* of the companion paper's branch-(N) reading. With transport valued in D_7 and gauge the orientation-preserving local frame relabellings ($\langle s \rangle$) at each hub, so that orientation is well-defined as in Gate-2), the gauge orbit of a *single face's* transport collapses to its orientation alone: every face is gauge-equivalent to e (orientation-preserving) or to r (orientation-reversing), and the \mathbb{Z}_7 position is entirely gauge. Hence a single face carries **no** gauge-invariant closure content. Consequently:

- **Pairwise compatibility is exhausted by orientation.** The triple-overlap programme's Theorem 3.1 ("every pairwise obstruction is already a Gate-2 obstruction") is **vindicated by the construction**, not falsified — the reverse of the companion paper's branch-(N) guess that κ is a pairwise obstruction.
- **κ is necessarily a loop holonomy.** The first gauge-invariant closure content is the \mathbb{Z}_7 -valued holonomy around closed loops. Through orientation-reversing faces the \mathbb{Z}_7 fibre inverts ($r s^k r^{-1} = s^{-k}$), so the natural home of κ — once the rotation connection is flat (the open closure-H-OC₇ condition of §4) — is the twisted cohomology

$$\kappa \in H^1(\Gamma_{\text{vac}}; \tilde{\mathbb{Z}}_7), \tilde{\mathbb{Z}}_7 = \mathbb{Z}_7 \text{ with the orientation } \mathbb{Z}_2 \text{ acting by inversion (monodromy} = [\text{or}]).$$

After Gate-2 ($[\text{or}]_{\text{vac}} = 0$) the local system trivializes on the vacuum and $\kappa \in H^1(\Gamma_{\text{vac}}; \mathbb{Z}_7)$, honest coefficients. The H^1 membership is conditional on per-plaquette flatness (Corollary 3.3); absent it, the holonomy is $\tilde{\mathbb{Z}}_7$ -valued without descending to a class.

The refined obstruction tower (Corollary 3.3, [Proven; H^1 membership conditional on flatness]). Pairwise / single-face: orientation only, killed by Gate-2. Degree 1 (H^1 , under closure- $H\text{-OC}_7$ flatness): $\kappa \in H^1(\Gamma_{\text{vac}}; \mathbb{Z}_7)$, the **lowest** Gate-3 obstruction, with minimal support a 3-hub loop — reconciling with the triple-overlap programme's " $n \geq 3$ realizers." Degree 2 (H^2): the orientation assembly class and the \mathbb{Z}_7 assembly class. Both obstruction families sit at or above the loop level; none is pairwise.

Three kernels and the reframed question (§5). The construction lets the central question sharpen from "does the kernel exist?" to "does any kernel component *survive*?" Three notions separate: the algebraic kernel $N_{\text{alg}} = \ker(\pi_{\text{or}})$ ($\cong \mathbb{Z}_7$ under the strong model); the holonomy kernel $N_{\text{hol}} \subseteq N_{\text{alg}}$, detectable only through loops; and the observable kernel $N_{\text{obs}} \subseteq N_{\text{alg}}$, surviving the operational quotient. Theorem 3.1 forces $N_{\text{obs}} \subseteq N_{\text{hol}} \subseteq N_{\text{alg}}$ — single-face purity is exactly the statement that observable kernel content must be holonomy content. Because \mathbb{Z}_7 is simple, each level is 0 or \mathbb{Z}_7 entire. The surviving invariant is therefore a residual cohomology class $\kappa \in H^1(\Gamma_{\text{vac}}; \mathbb{N})$, not a local overlap datum, and the programme reduces to one fork: **A**, $N_{\text{obs}} = 0$ (the kernel is operationally redundant, Gate-3 has one obstruction); or **B**, $N_{\text{obs}} = \mathbb{Z}_7$ (a residual closure-holonomy sector survives, Gate-3 has two). Which value N_{obs} takes — the operational-quotient step — is open.

What remains open, now sharply (§4). The verdict between rigidity ($\kappa \equiv 0$ forced by admissibility) and surplus ($\kappa \neq 0$ admissibly realized) is the *occupancy* question, now precise: does admissibility permit nonzero $\kappa \in H^1(\Gamma_{\text{vac}}; \mathbb{Z}_7)$? This has the architecture of Gate-2 one notch up — a per-plaquette \mathbb{Z}_7 -flatness condition (the closure analogue of $H\text{-OC}$), then a class on H_1 — and it requires the action of the hub-completion, conservation, and competition rules on the rotation labels, which the primitives given here do not yet fix. We construct the group, locate the surplus, and pose the occupancy question; we do not close it.

Labels: [Proven], [Proven under the strong model], [Methodological], [Open: admissibility action on closure labels].

Contents

1. The Primitives, Stated as Carried Structure
 2. The Transport Group: D_7 , and Why Not \mathbb{Z}_{14} or $\text{AGL}(1,7)$
 3. Single Faces Are Pure Gauge: the Surplus Lives on Loops
 4. The Occupancy Question, Now Well-Posed
 5. Three Notions of Transport Kernel, and What the Question Has Become
 6. Limitations
 7. Conclusion
 8. Relationship to the Locality and Operational-Quotient Results
 9. Future Directions: The Kernel Fork
-

1. The Primitives, Stated as Carried Structure

We take the closure-frame primitives as given and restate them as the data on which transport acts.

Primitive 1.1 (hub). A hub is a local $K = 7$ closure cell: six boundary channels that close through a seventh, completing channel. The seven channels carry a *cyclic-successor order* — the order in which the boundary relations return through the completing state — with labels in \mathbb{Z}_7 . The hub additionally carries an orientation (read-direction) sign recording the sense in which the cyclic order is traversed. We call the pair (cyclic order, read-direction) the hub's **closure frame**.

Primitive 1.2 (shared face). A shared face is the interface along which two adjacent hubs identify part of their boundary data, carrying enough information to compare the two closure frames.

Primitive 1.3 (transport). Transport across a face is the re-expression of one hub's closure frame in the adjacent hub's frame — a frame isomorphism. **Orientation** is the parity (read-direction) of that isomorphism: orientation-preserving if the cyclic sense agrees, orientation-reversing if it is reversed.

Primitive 1.4 (the carried object). The object carried across a face is the sevenfold cyclic closure frame. Transport is therefore valued in the symmetry group of a sevenfold cyclic-successor structure, computed in §2.

The single modelling choice — and the fork the companion paper isolated — is whether transport carries the frame (Primitive 1.4, the *strong* reading) or forgets the closure labels and records only the boundary orientation (the *weak* reading). We compute both in §2 and carry both forward, noting that the strong reading is the one consonant with calling the construction a sevenfold closure.

2. The Transport Group: D_7 , and Why Not \mathbb{Z}_{14} or $AGL(1,7)$

Definition 2.1 (cyclic-successor frame). A sevenfold cyclic-successor frame is a set of seven labels carrying (i) a cyclic successor relation $x \mapsto x + 1 \pmod{7}$ and (ii) a chosen sense. Its automorphisms are the bijections preserving the successor relation up to sense.

Theorem 2.2 (the structure group is D_7) [Proven under the strong model]. Under Primitive 1.4, the transport group is the automorphism group of a sevenfold cyclic-successor frame,

$G \cong D_7 = \langle s, r \mid s^7 = r^2 = e, r s r^{-1} = s^{-1} \rangle, |G| = 14,$

with $\langle s \rangle \cong \mathbb{Z}_7$ the rotations (preserving the cyclic sense, $x \mapsto x + c$) and the seven reflections $r s^k$ reversing it ($x \mapsto -x + c$).

Proof. A map $\mathbb{Z}_7 \rightarrow \mathbb{Z}_7$ preserves the successor relation $x \mapsto x + 1$ in the orientation-preserving sense iff it is a translation $x \mapsto x + c$: indeed $(x + 1) + c = (x + c) + 1$, and any successor-preserving bijection is determined by the image of one label and the requirement that successors map to successors, forcing constant offset. It reverses the successor relation (orientation-reversing sense) iff it is $x \mapsto -x + c$: then $x + 1 \mapsto -(x + 1) + c = (-x + c) - 1$, the predecessor, which is succession in the reversed sense. The translations form \mathbb{Z}_7 ; adjoining one reversal gives the inversion automorphism, and the two together generate D_7 with the stated presentation. No other bijection preserves the cyclic-successor structure (up to sense), so the automorphism group is exactly D_7 .

Definition 2.3 (orientation projection) and its kernel [Proven under the strong model].

Since $\langle s \rangle$ is the index-2 rotation subgroup, the rotation/reflection parity

$$\pi_{\text{or}}: D_7 \rightarrow \mathbb{Z}_2, s^k \mapsto 0, r s^k \mapsto 1$$

is a well-defined surjective homomorphism (its kernel is a subgroup of index 2, and $\langle s \rangle$ is normal). Its kernel is the rotation subgroup,

$$\ker(\pi_{\text{or}}) = \langle s \rangle \cong \mathbb{Z}_7,$$

which is the closure channel of the companion paper. The orientation projection of Primitive 1.3 — preserving versus reversing the cyclic sense — coincides with this parity, since preserving the sense is exactly being a rotation.

Remark 2.4 (why not $\text{AGL}(1, 7)$) [Methodological]. The full affine group of \mathbb{Z}_7 , $\text{AGL}(1, 7) = \mathbb{Z}_7 \rtimes \text{Aut}(\mathbb{Z}_7) = \mathbb{Z}_7 \rtimes \mathbb{Z}_6$ of order 42, is the symmetry group of \mathbb{Z}_7 as an *abstract group* (or as an unordered affine line). It is too large here: a dilation $x \mapsto ax$ with a $a \notin \{1, -1\}$ sends the successor $x + 1$ to $ax + a$, which is a step of size a , not 1 — it does not preserve the cyclic *adjacency* of the closure. The closure frame is a cyclic-successor order, not merely a \mathbb{Z}_7 -set, so only $a = \pm 1$ survive, and $\text{AGL}(1, 7)$ collapses to $D_7 = \mathbb{Z}_7 \rtimes \{\pm 1\}$. Carrying the seven *as a closure cycle* is strictly more rigid than carrying it as a bare \mathbb{Z}_7 label, and this rigidity is what fixes $G = D_7$.

Theorem 2.5 (D_7 is forced over \mathbb{Z}_{14}) [Proven under the strong model]. Under Primitives 1.1–1.4, the orientation \mathbb{Z}_2 acts on the closure channel \mathbb{Z}_7 by inversion; hence $G \cong D_7$, not $G \cong \mathbb{Z}_{14} = \mathbb{Z}_7 \times \mathbb{Z}_2$.

Proof. Orientation is the read-direction of the closure frame (Primitive 1.3), and the read-direction of a cyclic-successor order is its sense. An orientation-reversing transport therefore reverses the cyclic sense, i.e. sends the successor relation to its opposite, i.e. acts on the labels by $x \mapsto -x$ — inversion of $\langle s \rangle$. The semidirect action of the orientation \mathbb{Z}_2 on \mathbb{Z}_7 is thus inversion, the nontrivial element of $\text{Aut}(\mathbb{Z}_7) \cong \mathbb{Z}_6$ of order 2, giving the dihedral extension D_7 . The direct

product \mathbb{Z}_{14} corresponds to the *trivial* action, in which the orientation generator commutes with $\langle s \rangle$ and so reverses read-direction while fixing the cyclic labels. But fixing the labels means preserving the cyclic sense, contradicting "reverses read-direction." Hence the trivial action is incompatible with the identification of orientation with cyclic sense, and \mathbb{Z}_{14} is excluded under the strong model.

Proposition 2.6 (the modelling fork) [Proven]. Two readings of Primitive 1.3 give two transport groups.

- **Weak reading** (transport records only induced boundary orientation, forgetting the closure labels). Then transport takes values in the orientation \mathbb{Z}_2 alone:

$$G \cong \mathbb{Z}_2, \ker(\pi_{\text{or}}) = \{1\},$$

so the closure channel is empty and rigidity holds outright (companion paper, Proposition 6.2). The cost is that the construction's defining seven plays no role in transport — it is substrate-internal but not transported.

- **Strong reading** (transport carries the closure frame, Primitive 1.4). Then by Theorems 2.2 and 2.5,

$$G \cong D_7, \ker(\pi_{\text{or}}) = \mathbb{Z}_7, 7 \mid |G|,$$

so the closure channel is structurally nontrivial and, by the companion paper's Proposition 6.3, the seven is forced into the kernel.

Proof. Immediate from the two readings and §2's results. The naturalness judgement — that a closure construction should transport its closure — is a modelling stance, not a theorem, and we mark it as such: the strong reading is the one under which the seven is not idle, but the weak reading is logically available and yields rigidity trivially.

We proceed under the strong reading, flagging at each consequence that it is conditional on it.

3. Single Faces Are Pure Gauge: the Surplus Lives on Loops

The gauge group is the local frame relabellings at each hub. A relabelling of a hub's closure frame is an element of the structure group acting on that hub's frame; to keep orientation well-defined as a \mathbb{Z}_2 cochain (as Gate-2 requires), we take the gauge at each hub to be the orientation-preserving relabellings, the rotations $\langle s \rangle \cong \mathbb{Z}_7$. Transport on a face f between hubs α, β transforms under gauge $\gamma_{\alpha}, \gamma_{\beta} \in \langle s \rangle$ in the standard transition-function way,

$$g_f \mapsto \gamma_\beta g_f \gamma_\alpha^{-1}.$$

Theorem 3.1 (a single face carries no gauge-invariant closure content) [Proven under the strong model]. Under rotation gauge, every admissible single-face transport $g_f \in D_7$ is gauge-equivalent to a representative determined by its orientation alone: to e if g_f is orientation-preserving, to r if g_f is orientation-reversing. The closure (\mathbb{Z}_7) position of a single face is entirely gauge.

Proof. If $g_f = s^k$ (orientation-preserving), then for $\gamma_\alpha = s^a$, $\gamma_\beta = s^b$,

$$\gamma_\beta g_f \gamma_\alpha^{-1} = s^b s^k s^{-a} = s^{k+b-a},$$

which is e for the choice $b-a = -k$; the orbit is all of $\langle s \rangle$, with the identity in it. If $g_f = r s^k$ (orientation-reversing), then using $s^b r = r s^{-b}$,

$$\gamma_\beta g_f \gamma_\alpha^{-1} = s^b (r s^k) s^{-a} = r s^{-b} s^k s^{-a} = r s^{k-a-b},$$

which is r for $k-a-b = 0$; the orbit is all seven reflections, with r in it. In both cases the gauge orbit is an entire π_{or} -fibre, so the only gauge-invariant of a single face is $\pi_{\text{or}}(g_f) = \text{its orientation}$. The \mathbb{Z}_7 content is gauge.

Corollary 3.2 (pairwise is exhausted by orientation; the triple-overlap premise is vindicated) [Proven under the strong model]. Since a single face's only gauge-invariant is its orientation, every pairwise (single-face) obstruction is an orientation obstruction, hence already a Gate-2 obstruction. The triple-overlap programme's Theorem 3.1 ("every pairwise obstruction is already a Gate-2 obstruction") holds under the strong model — and holds *whether or not* the seven is transported, since the \mathbb{Z}_7 content, when present, is pairwise-trivial by Theorem 3.1.

Proof. Immediate from Theorem 3.1: orientation is the entire pairwise content.

This **corrects the level of the companion paper's branch-(N) reading**. That paper, working abstractly, allowed κ to be a *pairwise* obstruction "below the triple level," which would have falsified the triple-overlap thesis. The explicit primitives show the surplus cannot be pairwise: single faces are pure gauge, so κ lives only on loops, and the triple-overlap localization is therefore correct rather than undercut. The companion paper's Theorem 4.1 (the reduction $\kappa = 0 \Leftrightarrow \text{H-R} \Leftrightarrow \text{triple-overlap premise}$) is unaffected as a logical equivalence; what the construction corrects is only the *level* at which κ can be nonzero — never pairwise, only on loops — and in doing so it vindicates the triple-overlap premise rather than overturning anything load-bearing.

Corollary 3.3 (the closure holonomy is $\tilde{\mathbb{Z}}_7$ -valued; its descent to \mathbf{H}^1 is conditional on flatness) [Proven; \mathbf{H}^1 membership Conditional]. The first gauge-invariant closure content is the holonomy around closed loops. For a loop $\ell = f_n \circ \dots \circ f_1$ the holonomy $H(\ell) = g_{\{f_n\}} \dots g_{\{f_1\}} \in D_7$ is gauge-invariant up to basepoint conjugation; its orientation part $\pi_{\text{or}}(H(\ell))$ is the orientation holonomy ([or] paired with ℓ). On orientation-trivial loops $H(\ell) \in \mathbb{Z}_7$, and since orientation-reversing faces conjugate \mathbb{Z}_7 by inversion ($r s^k r^{-1} = s^{-k}$), the holonomy is valued in the local system $\tilde{\mathbb{Z}}_7$ (\mathbb{Z}_7 twisted by the orientation \mathbb{Z}_2 -action, monodromy [or]):

(proven, no flatness needed) $H(\ell) \in \tilde{\mathbb{Z}}_7$ for every orientation-trivial loop ℓ .

For this loop holonomy to descend to a cohomology class — to depend only on the homology class of ℓ — the rotation transport must be a *flat* \mathbb{Z}_7 connection: trivial rotation holonomy around every elementary plaquette (the closure analogue of H-OC, "closure H-OC", §4 step 1). Under that condition, and only then,

(conditional on per-plaquette flatness) $\kappa \in H^1(\Gamma_{\text{vac}}; \tilde{\mathbb{Z}}_7)$, untwisted to $H^1(\Gamma_{\text{vac}}; \mathbb{Z}_7)$ after Gate-2 ([or]_{vac} = 0).

The flatness is precisely the open condition deferred to §4; we therefore do not claim $\kappa \in H^1$ as established here, only that the holonomy is $\tilde{\mathbb{Z}}_7$ -valued and that it descends to H^1 exactly under closure H-OC₇.

The obstruction tower (under flatness, so that the H^1 classes exist) is then:

- **pairwise / single face:** orientation only — killed by Gate-2 (Corollary 3.2);
- **degree 1, H^1 :** $\kappa \in H^1(\Gamma_{\text{vac}}; \mathbb{Z}_7)$, the lowest Gate-3 obstruction, minimal support a 3-hub loop;
- **degree 2, H^2 :** the orientation assembly class and the \mathbb{Z}_7 assembly class.

Proof. The $\tilde{\mathbb{Z}}_7$ -valued statement: products of transition functions around a loop are conjugation-invariant (gauge-invariant up to basepoint), the orientation projection is the \mathbb{Z}_2 holonomy, and on orientation-trivial loops the holonomy lies in the kernel \mathbb{Z}_7 , twisted by the conjugation action of orientation-reversing steps, which is inversion ($r s^k r^{-1} = s^{-k}$) — none of this uses flatness. The descent to H^1 : a connection's holonomy factors through homology iff it is flat (trivial around every contractible loop, equivalently every elementary plaquette); this is the closure H-OC₇ condition, which is open (§4). The minimal nontrivial 1-cycle is a loop through three distinct hubs (two hubs give a single face, which is pure gauge by Theorem 3.1), so κ 's minimal support has three realizers, matching the triple-overlap programme's $n \geq 3$.

Remark 3.3a (Γ_{vac} carries 2-cells, so flatness is a genuine condition) [Methodological].

The flatness condition is non-vacuous precisely because Γ_{vac} is not a pure 1-complex: §4's plaquette and bounding-loop language presupposes elementary 2-cells, around which the rotation holonomy may or may not be trivial. Were Γ_{vac} a bare 1-complex (no 2-cells), $H^1 = \text{Hom}(\pi_1, \mathbb{Z}_7)$ would be automatic and there would be no flatness condition — but then §4's plaquette analysis would be empty. Since §4 has plaquettes, Γ_{vac} carries 2-cells, flatness is a real per-plaquette constraint, and $\kappa \in H^1$ is genuinely conditional on it.

Remark 3.4 (why the level matters) [Methodological]. The companion paper posed H-R as a per-face condition; Theorem 3.1 shows that at the per-face level rigidity is *automatic*, so the per-face statement, while true, is vacuous and locates nothing. The content of rigidity is therefore entirely the loop statement: H-R holds iff the \mathbb{Z}_7 loop holonomy κ vanishes for every admissible vacuum configuration. This is not a weakening of the companion paper's logic but a correction of its level — the same kind of relocation the matching-cochain paper made when it moved the orientation residue from bounding loops (automatically trivial) to non-bounding loops (where the

content lives). Here the closure surplus moves from single faces (automatically trivial) to loops (where κ lives).

4. The Occupancy Question, Now Well-Posed

The construction settles existence and locates the surplus; it does not settle occupancy. Under the strong model the closure channel is \mathbb{Z}_7 (Proposition 2.6), and κ descends to a class in $H^1(\Gamma_{\text{vac}}; \mathbb{Z}_7)$ once the rotation connection is flat (Corollary 3.3, step 1 below). The rigidity-versus-surplus verdict is now exactly:

Does admissibility permit a nonzero $\kappa \in H^1(\Gamma_{\text{vac}}; \mathbb{Z}_7)$, or force $\kappa = 0$ on every admissible vacuum configuration?

Nonzero, admissibly realized \rightarrow branch (N): a genuine \mathbb{Z}_7 closure-holonomy sector, the discovery outcome. Forced to zero \rightarrow branch (R): rigidity by non-occupancy (the companion paper's B1, now the only surviving route to rigidity under the strong model, since Theorem 2.5 removed the entanglement route).

The architecture of the occupancy check [Methodological]. This question has the same shape as Gate-2, one notch up in coefficients. Gate-2 analysed an orientation (\mathbb{Z}_2) cochain: a per-plaquette closure condition (H-OC, the matching-cochain paper) made it a cocycle, after which its class lived on H_1 and the pass condition was vanishing of that class. The closure channel asks the identical questions with \mathbb{Z}_7 in place of \mathbb{Z}_2 :

1. **Closure (flatness) of the \mathbb{Z}_7 connection.** Is the product of the rotation parts of transport trivial around every elementary plaquette — the \mathbb{Z}_7 analogue of H-OC? If so, the rotation transport is a flat \mathbb{Z}_7 connection and κ is a well-defined class in $H^1(\Gamma_{\text{vac}}; \mathbb{Z}_7)$.
2. **The class on H_1 .** Given flatness, κ is determined by its values on the non-bounding loops of Γ_{vac} , and can be nonzero iff $H^1(\Gamma_{\text{vac}}; \mathbb{Z}_7) \neq 0$ and admissibility does not force the class to vanish.
3. **Admissible forcing.** Whether the hub-completion, conservation, and competition rules constrain the \mathbb{Z}_7 holonomy to be trivial even on non-bounding loops — the genuine substrate question.

What is needed to close it [Open: admissibility action on closure labels]. The single missing input is the action of admissibility on the *rotation* labels — how hub completion, conservation, and competition constrain the \mathbb{Z}_7 part of transport around a plaquette and around a non-bounding loop. The primitives of §1 fix the carried object (the closure frame) and hence G , π_{or} , and the location of κ ; they do not fix how the substrate's completion rules act on the seven. That action is the analogue, for the closure channel, of the orientation-coherence data the matching-cochain programme used for the \mathbb{Z}_2 channel, and it is the next primitive to write down.

We therefore hand forward a sharply posed problem rather than a verdict: construct the \mathbb{Z}_7 connection from the admissibility rules, test per-plaquette flatness (closure H-OC₇), and compute the resulting class in $H^1(\Gamma_{\text{vac}}; \mathbb{Z}_7)$. The obstruction count of Gate-3 is one (orientation assembly only) if that class is forced to vanish, and two (plus the \mathbb{Z}_7 κ -class) if it is not.

5. Three Notions of Transport Kernel, and What the Question Has Become

The construction lets the central question be stated more sharply than "does the transport group have a nontrivial kernel under π_{or} ?" Theorem 3.1 shows that a nontrivial *algebraic* kernel need not be *observable*: the kernel can be structurally present (\mathbb{Z}_7) yet carry nothing detectable on a single face. The right question is therefore not existence of the kernel but **survival** of any kernel component through gauge identification, loop-detection, and the operational quotient. Three notions of kernel separate these, and the paper's results arrange them into a hierarchy.

Definition 5.1 (three transport kernels). Let $N = \ker(\pi_{\text{or}})$ be the closure channel. Define:

1. **Algebraic kernel** $N_{\text{alg}} = N$ — the full set of transport elements acting trivially on orientation. (Under the strong model, $N_{\text{alg}} \cong \mathbb{Z}_7$.)
2. **Holonomy kernel** $N_{\text{hol}} \subseteq N_{\text{alg}}$ — the kernel content detectable only through nontrivial transport cycles: the part of N that appears in some loop holonomy and is not removable by gauge.
3. **Observable kernel** $N_{\text{obs}} \subseteq N_{\text{alg}}$ — the kernel content producing operationally distinguishable physical effects, after the operational quotient is applied.

These are stated as subobjects of N_{alg} ; Proposition 5.5 will present N_{obs} equivalently as a quotient, $N_{\text{obs}} = N_{\text{hol}} / \sim_{\text{op}}$. The two readings — subgroup and quotient — are different categorical operations in general, but they coincide here: in the $\{0, \mathbb{Z}_7\}$ regime forced by the simplicity of \mathbb{Z}_7 (Theorem 5.3), every subgroup is also the image of a quotient and conversely, so N_{obs} may be read either way without ambiguity. We use the two interchangeably and note the coincidence is exactly what the simplicity buys.

Proposition 5.2 (the kernel hierarchy) [Proven under the strong model, for the first inclusion; $N_{\text{obs}} \subseteq N_{\text{hol}}$ by Theorem 3.1].

$$N_{\text{obs}} \subseteq N_{\text{hol}} \subseteq N_{\text{alg}},$$

with each inclusion powered by a distinct mechanism:

- $N_{\text{hol}} \subseteq N_{\text{alg}}$ holds by definition: holonomy-detectable content is kernel content, since a loop with trivial orientation holonomy carries holonomy valued in N (Corollary 3.3).

- $N_{\text{obs}} \subseteq N_{\text{hol}}$ holds by **Theorem 3.1**: a single face carries no gauge-invariant closure content, so no element of the kernel is observable below the loop level; any observable kernel content must therefore be holonomy content. This inclusion is not a definitional convenience — it is the structural consequence of single-face purity, and it is the precise sense in which "the surplus lives on loops."

Proof. $N_{\text{hol}} \subseteq N_{\text{alg}}$: immediate from the definitions and Corollary 3.3, since holonomy on orientation-trivial loops is N -valued. $N_{\text{obs}} \subseteq N_{\text{hol}}$: by Theorem 3.1 every single-face transport is gauge-equivalent to e or r , so its closure (N) part is gauge-trivial; an operationally distinguishable effect of kernel content cannot then arise on a single face, and must be carried by a loop, i.e. lie in N_{hol} .

Theorem 5.3 (Kernel Reduction Theorem) [Proven under the strong model]. Under the strong model the algebraic kernel is completely determined,

$$N_{\text{alg}} = \ker(\pi_{\text{or}}) \cong \mathbb{Z}_7,$$

and the hierarchy

$$N_{\text{obs}} \subseteq N_{\text{hol}} \subseteq N_{\text{alg}} \cong \mathbb{Z}_7$$

has each level fixed by a strictly later stage of the analysis:

1. N_{alg} is fixed by transport geometry alone (the carried object; §2).
2. N_{hol} is fixed by loop transport (gauge plus holonomy; Theorem 3.1, Corollary 3.3).
3. N_{obs} is fixed only after admissibility and the operational quotient (§4, and 5.5 below).

Consequently the only remaining uncertainty lies in the two reductions

$$N_{\text{alg}} \rightarrow N_{\text{hol}} \rightarrow N_{\text{obs}},$$

and — because $N_{\text{alg}} \cong \mathbb{Z}_7$ is **simple** — each reduction is all-or-nothing as a subgroup: $N_{\text{hol}} \in \{0, \mathbb{Z}_7\}$ and $N_{\text{obs}} \in \{0, \mathbb{Z}_7\}$. The closure channel survives each stage entire or not at all; there is no proper nontrivial residue.

Proof. The inclusions and their mechanisms are Proposition 5.2; $N_{\text{alg}} \cong \mathbb{Z}_7$ is Definition 2.3. The subgroup statement is the simplicity of \mathbb{Z}_7 (prime order, hence no proper nontrivial subgroups): any subgroup, in particular N_{hol} and N_{obs} , is 0 or all of \mathbb{Z}_7 .

The conceptual content. The Reduction Theorem is the point at which the transport problem changes type. Once N_{alg} is determined, identifying the transport group is finished business — there is no further group to find. What remains is entirely the question of which of the two reductions $N_{\text{alg}} \rightarrow N_{\text{hol}}$ and $N_{\text{hol}} \rightarrow N_{\text{obs}}$ are trivial. The transport problem has therefore **ceased to be a group-identification problem and become a kernel-reduction problem**. This is why an analysis that begins by asking "what is G ?" gives way to analyses that ask about admissibility, loop structure, operational quotients, and locality: those are not new subjects but

the two reductions of this theorem, attacked from different directions. Each later programme is a determination of one arrow in $N_alg \rightarrow N_hol \rightarrow N_obs$.

Remark 5.4 (simplicity governs the channel, not the class) [Methodological]. Theorem 5.3's all-or-nothing statement is about the kernel as a *coefficient channel* — whether \mathbb{Z}_7 survives a reduction stage as a group. It must not be misread as a claim about the cohomology class κ . The simplicity of \mathbb{Z}_7 implies that survival of the coefficient channel is all-or-nothing; it does **not** imply that κ is binary. Once the coefficient channel survives ($N_obs = \mathbb{Z}_7$), κ retains the full structure of $H^1(\Gamma_vac; \mathbb{Z}_7)$, which can contain many nonzero classes — the class is binary in *whether the channel carries it*, not in *how much* it carries. Channel survival (a subgroup question, two-valued) and class content (a cohomology question, generally large) are distinct, and only the first is governed by simplicity.

Proposition 5.5 (operational-quotient interpretation) [Methodological / Definitional]. A distinction with no operational witness is not physical; the operational programme removes kernel content lacking such a witness by quotienting. Writing \sim_op for operational indistinguishability, the observable kernel is the operational quotient of the holonomy kernel,

$$N_obs = N_hol / \sim_op.$$

The existence of a nontrivial *algebraic* kernel ($N_alg \cong \mathbb{Z}_7$, Theorem 5.3) therefore does **not** imply a nontrivial physical sector: it implies one only if N_hol is nonzero (loop survival, the occupancy question of §4) *and* that holonomy carries an operational witness (survival of the operational quotient). By the simplicity of \mathbb{Z}_7 , both are all-or-nothing. The central open question is thus sharpened from "does κ exist?" to:

Does any nontrivial class in $H^1(\Gamma_vac; \mathbb{Z}_7)$ survive the operational quotient as an operationally distinguishable invariant?

Proof. The equality is the definition of N_obs as the operationally distinguishable part of N_hol ; the implication and the all-or-nothing follow from Theorem 5.3 and Definition 5.1.

Remark 5.6 (the surviving inclusion is itself the open question) [Methodological]. Whether $N_obs = N_hol$ or $N_obs = 0$ (the only two options, by simplicity) is **open**, and is exactly the operational-quotient content of the occupancy question (§4). It is conceivable that loop holonomy is gauge-invariant ($N_hol = \mathbb{Z}_7$) yet operationally indistinguishable ($N_obs = 0$) once the operational quotient is applied — admissibility and gauge equivalence may remove kernel content before it becomes physically observable. The programme's stages cut the chain at successively later points: the matching-cochain analysis fixed the orientation (\mathbb{Z}_2) layer; the present construction fixes $N_alg = \mathbb{Z}_7$ and shows closure content survives gauge only as loop holonomy ($N_obs \subseteq N_hol$, Theorem 3.1); the remaining locality and operational-quotient questions decide where N_hol and N_obs land in $\{0, \mathbb{Z}_7\}$. We state the fork as genuinely balanced: nothing in the chain to date establishes $N_obs = 0$, and nothing establishes $N_obs = \mathbb{Z}_7$.

Reframing κ as a residual cohomology class [Methodological]. In the earliest framing κ was "possible hidden transport content," a local degree of freedom on overlaps. The construction

supersedes that reading: the natural and gauge-correct home of κ is a residual cohomology class valued in the transport kernel,

$$\kappa \in H^1(\Gamma_{\text{vac}}; \tilde{N}), N = \ker(\pi_{\text{or}}),$$

with \tilde{N} the local system of N twisted by the orientation action (Corollary 3.3; under the strong model $N = \mathbb{Z}_7$, $\tilde{N} = \tilde{\mathbb{Z}}_7$, and after Gate-2 the coefficients are untwisted). κ is therefore not a local overlap degree of freedom but a cohomology class of residual transport — an element of N_{hol} , observable only if it survives into N_{obs} . This is the model-independent statement: whatever N turns out to be, the surviving closure invariant is a class in $H^1(\Gamma_{\text{vac}}; N)$, not a face datum.

The fork, stated cleanly. The construction and the hierarchy reduce the entire question above Gate-2 to a choice between two endpoints.

- **Fork A — rigidity (redundancy).** $N_{\text{obs}} = 0$: the kernel is operationally redundant, removed by admissibility and gauge before any observable survives. Then the orientation quotient $G / N \cong \mathbb{Z}_2$ carries all physical transport content, the triple-overlap localization stands with a single (orientation-assembly) obstruction, and Gate-3 carries one obstruction.
- **Fork B — residual holonomy.** $N_{\text{obs}} \neq 0$: a kernel component survives as a nonlocal loop observable. Then $\kappa \in H^1(\Gamma_{\text{vac}}; N)$ is a genuine surviving transport invariant, Gate-3 carries two obstruction families, and a residual closure-holonomy sector has been discovered.

Deciding between Fork A and Fork B is the operational-quotient question of Remark 5.6, and it needs the admissibility action on the kernel labels (§4) together with the operational quotient. The construction has reduced the programme to this fork; it does not decide it.

6. Limitations

Conditional on the strong model. All of §2.2 onward — $G \cong D_7$, $\ker(\pi_{\text{or}}) = \mathbb{Z}_7$, the loop-holonomy location of κ — holds under the strong reading of Primitive 1.3 (transport carries the closure frame). Under the weak reading $G \cong \mathbb{Z}_2$, the channel is empty, and rigidity is immediate (Proposition 2.6). The dependency's reach is worth stating plainly: the central physical question (Fork A versus B) is *nontrivial only under the strong model* — the weak model yields Fork A trivially, an empty channel and rigidity outright — so the unproven modelling choice is load-bearing not merely for which branch holds but for the occupancy question existing at all. The strong reading is supported only by the naturalness stance that a closure construction should transport its closure (the "idle seven" argument); it is not a theorem, and we mark every strong-model result accordingly. [Proven under the strong model.]

The occupancy verdict is not reached. The construction proves the channel exists and is \mathbb{Z}_7 , locates κ in $H^1(\Gamma_{\text{vac}}; \mathbb{Z}_7)$, and poses the occupancy question; it does not answer it, because the

action of admissibility on the rotation labels is not among the primitives supplied. Whether κ can be nonzero is [Open: admissibility action on closure labels].

The orientation assembly class is still unanalysed. The degree-2 orientation class of the companion paper (its Theorem 5.1(R)(c), the $H^2(\Gamma_{\text{vac}}; \mathbb{Z}_2)$ assembly obstruction) is named here only to place it in the tower (Corollary 3.3); whether it is nonzero is a separate \mathbb{Z}_2 computation, untouched.

Gauge taken orientation-preserving. Theorem 3.1 uses rotation gauge $\langle s \rangle$, the orientation-preserving relabellings, so that orientation is a well-defined cochain (as Gate-2 requires). If the gauge group were taken to include reflections, even orientation would fail to be a single-face invariant; the restriction to rotation gauge is what makes the orientation cochain meaningful, and it is the same restriction implicit in the Gate-2 construction.

The framework is taken as given. The closure-frame primitives, the Gate-2 verdict, the matching-cochain identification, and the companion paper's reduction and reconciliation are taken as established within the VERSF / $K = 7$ framework. The group theory of §2–§3 is proved here from the primitives and standard gauge-transition facts, and depends on no prior paper for its steps. "Established," where used, means within that framework and under the stated model.

7. Conclusion

Given the closure-frame primitives, the transport group is no longer a free unknown. Transport that carries the sevenfold closure frame is valued in the symmetry group of a cyclic-successor heptagon — the dihedral group D_7 of order fourteen, with the seven rotations as the orientation-preserving closure channel and the seven reflections reversing the cyclic sense. The orientation projection is the rotation/reflection parity, its kernel is exactly \mathbb{Z}_7 , and the dihedral form is forced over the abelian \mathbb{Z}_{14} because orientation here is the read-direction and the read-direction of a closure cycle is its sense — reversing one reverses the other, which is inversion. The larger affine group is excluded because a closure *cycle* is more rigid than a bare \mathbb{Z}_7 label: only ± 1 preserve its adjacency. The sole alternative is the weak reading in which transport forgets the closure and records only a sign, giving $G \cong \mathbb{Z}_2$ and rigidity outright — at the price of making the construction's seven idle in transport.

The principal finding corrects the *level* of the companion paper's branch-(N) reading. One expects, abstractly, that a populated closure channel would show itself when two hubs meet across a face. It does not: a single face has enough relabelling freedom at its two ends to absorb any sevenfold twist, leaving only orientation behind. Pairwise compatibility is therefore exhausted by orientation, the triple-overlap programme's claim that nothing survives pairwise is vindicated rather than overturned, and the sevenfold surplus — if admissibility produces it — is forced up to the loop level, where it lives as a \mathbb{Z}_7 holonomy coupled to orientation, a class in $H^1(\Gamma_{\text{vac}}; \mathbb{Z}_7)$ on the Gate-2-passed vacuum (conditional on the closure-H-OC₇ flatness of §4).

The lowest Gate-3 obstruction is this κ , on a 3-hub loop; above it sit the degree-2 assembly classes.

What is left is genuinely a single, well-posed question, with the architecture of Gate-2 raised one notch in its coefficients: does admissibility, acting on the rotation labels of the closure frame, force the \mathbb{Z}_7 holonomy to vanish on every admissible vacuum loop, or admit a nonzero class? The first answer is rigidity by non-occupancy and a one-obstruction Gate-3; the second is a genuine \mathbb{Z}_7 closure-holonomy sector and a two-obstruction Gate-3. Answering it needs one more primitive — how the substrate's completion, conservation, and competition rules act on the seven — which the data given here do not contain. We have built the group, forced its form, located the surplus, and posed the check; the closure connection and its holonomy are the next step.

Final observation: three questions, previously conflated, now separated. The programme had run together three questions that this paper, with the kernel hierarchy of §5, pulls apart:

1. **Does a transport kernel exist?** — Answered positively under the strong model: $\ker(\pi_{\text{or}}) \cong \mathbb{Z}_7$. N_{alg} is *identified*.
2. **Does any kernel content survive gauge?** — Answered by Theorem 3.1: single-face kernel content is pure gauge, so only loop holonomy can survive. N_{hol} is *localized* — to $H^1(\Gamma_{\text{vac}}; \mathbb{Z}_7)$, the loop level, and nowhere below it.
3. **Does any surviving kernel content remain operationally observable?** — **Open.** $N_{\text{obs}} \in \{0, \mathbb{Z}_7\}$ is undetermined; nothing in the chain to date forces either value.

Accordingly the transport programme has evolved from a problem of group *identification* to a problem of operational *survival*. The remaining task is not to determine whether the kernel exists — it does — but whether any kernel-valued cohomology class survives admissibility and the operational quotient as a physically distinguishable invariant. This is the joint conclusion toward which the separate threads converge — the transport-group construction, the one-residue-or-two reduction, the operationally-invisible-boundary analysis, the locality decision, and the transport-holonomy programme — each fixing one arrow of the reduction chain. Stated as plainly as the chain permits:

Reduction status. We have identified N_{alg} ($\cong \mathbb{Z}_7$, the algebraic kernel). We have localized N_{hol} (to loop holonomy, $H^1(\Gamma_{\text{vac}}; \mathbb{Z}_7)$, with nothing surviving below the loop level). The remaining problem is determining N_{obs} — whether the localized channel survives the operational quotient as a physically distinguishable invariant.

We make no claim as to which value N_{obs} takes. The fork is balanced: single-face redundancy and loop-localization are established; operational redundancy ($N_{\text{obs}} = 0$) is not, and neither is operational survival ($N_{\text{obs}} = \mathbb{Z}_7$).

In one line:

Transport that carries the $K = 7$ closure frame is valued in D_7 — seven rotations forming the closure channel $\mathbb{Z}_7 = \ker(\pi_{\text{or}})$, seven reflections carrying orientation — with the dihedral form forced because orientation is the cyclic read-direction and its reversal is

inversion, and the affine group excluded because a closure cycle admits only ± 1 ; but a single face turns out to be pure gauge in everything except orientation, so the sevenfold surplus is never pairwise — vindicating the triple-overlap localization — and lives instead as a \mathbb{Z}_7 loop holonomy $\kappa \in H^1(\Gamma_{\text{vac}}; \mathbb{Z}_7)$, the lowest Gate-3 obstruction, whose vanishing-or-not is the now-sharp occupancy question that a single further primitive — the admissibility action on the rotation labels — will decide.

8. Relationship to the Locality and Operational-Quotient Results

Stated in the language of the locality, operational-quotient, and transport-holonomy analyses, the results above take a cleaner form, and the genuinely open question separates from what is settled. The two statements below restate §3 and §5 in that form; they add interpretation, not new derivation.

8.1 The Single-Face Rigidity Principle

Principle 8.1 (single-face rigidity) [Proven under the strong model (Theorem 3.1); conjectured model-independent]. For any admissible single-overlap transport element g_f ,

$$g_f \sim e \text{ or } g_f \sim r,$$

modulo admissibility and gauge equivalence. Consequently no element of $\ker(\pi_{\text{or}})$ is observable at the level of a single overlap, and orientation is the only local transport observable; all additional transport structure is necessarily global (loop-supported).

Under the strong model this is exactly Theorem 3.1, established for $G \cong D_7$ with orientation-preserving (rotation) gauge. The reason it is *plausibly model-independent* — and so stated as a principle, conjectural beyond the strong model — is that the gauge argument of Theorem 3.1 uses only that gauge acts by $g_f \mapsto \gamma_\beta g_f \gamma_\alpha^{-1}$ with the γ orientation-preserving, and that the orbit fills the π_{or} -fibre: it does not use the dihedral structure specifically. Whatever the true carried object, if its gauge group acts transitively on each π_{or} -fibre of a single face, single-face rigidity follows. We mark the general statement conjectural because that transitivity is established here only for the cyclic-successor frame.

8.2 Reinterpretation of the κ Sector

The κ -sector was first introduced as the possibility that transport carries information beyond the orientation projection — implicitly, a local overlap degree of freedom. The construction supersedes that reading. By Principle 8.1 no kernel content is local, so κ cannot be an overlap coordinate; its gauge-correct home is a residual transport-holonomy class valued in the kernel,

$\kappa \in H^1(\Gamma_{\text{vac}}; N)$, $N = \ker(\pi_{\text{or}})$,

with the coefficients carrying the orientation twist off the vacuum and untwisted after Gate-2 (§5; Corollary 3.3). Under this reading local overlaps carry only orientation; kernel information survives only through closed transport cycles; and κ is a cohomological invariant, not a local transport coordinate. This dissolves the apparent tension between the rigidity programme and the holonomy programme: local transport may be fully rigid (Principle 8.1) while a nontrivial residual class survives globally ($\kappa \in H^1(\Gamma_{\text{vac}}; N)$). The two are not competitors but the local and global halves of one structure.

8.3 The Locality and Operational-Quotient Programmes as the Two Reductions

The Kernel Reduction Theorem (5.3) identifies exactly two reductions, $N_{\text{alg}} \rightarrow N_{\text{hol}}$ and $N_{\text{hol}} \rightarrow N_{\text{obs}}$, and the later programmes can be read as the determination of one each.

- The **locality programme** determines the first reduction: whether loop transport carries surviving closure content, i.e. whether

$N_{\text{hol}} = N_{\text{alg}} \cong \mathbb{Z}_7$ or $N_{\text{hol}} = 0$.

(By Theorem 5.3 these are the only two options — simplicity leaves no third.)

- The **operational-indistinguishability programme** determines the second reduction, *given* the first: whether surviving loop holonomy carries an operational witness, i.e. whether

$N_{\text{obs}} = N_{\text{hol}}$ or $N_{\text{obs}} = 0$.

These are not independent developments but successive — and conditional — stages of one kernel-reduction chain. The second stage is live only on the branch where the first returns $N_{\text{hol}} = \mathbb{Z}_7$: if locality returns $N_{\text{hol}} = 0$, the operational question is moot, $N_{\text{obs}} = 0$ is forced, and the chain terminates at the first arrow. The locality programme is therefore strictly upstream and can short-circuit the operational one; only if it finds loop survival does the operational quotient have anything to act on. Read this way, the locality, operational-quotient, and transport-holonomy analyses are coordinated attacks on the single chain $N_{\text{alg}} \rightarrow N_{\text{hol}} \rightarrow N_{\text{obs}}$, each resolving one arrow, in order.

9. Future Directions: The Kernel Fork

The remaining transport question reduces to a fork concerning the *physical* status of $N = \ker(\pi_{\text{or}})$ — not its algebraic existence, which §2 settles, but whether any part of it survives the operational quotient (§5, Definition 5.1).

Branch A — operational rigidity. The kernel is entirely gauge redundancy:

$$N_{\text{obs}} = 0.$$

Then $G / N \cong \mathbb{Z}_2$ carries all physically meaningful transport content, transport reduces completely to orientation, and Gate-3 carries one obstruction (the orientation assembly class alone).

Branch B — residual holonomy. Some kernel content survives the operational quotient:

$$N_{\text{obs}} \neq 0 \text{ (hence } N_{\text{hol}} \neq 0, \text{ since } N_{\text{obs}} \subseteq N_{\text{hol}}).$$

Then physical transport contains an additional cohomological sector

$$\kappa \in H^1(\Gamma_{\text{vac}}; N),$$

and Gate-3 carries two obstruction families. (By the simplicity of \mathbb{Z}_7 , survival is all-or-nothing at the channel level: N_{obs} is either 0 or \mathbb{Z}_7 , so Branch B asserts $N_{\text{obs}} = \mathbb{Z}_7$ — the closure channel survives entire as an observable coefficient group — together with a nonzero class κ . There is no proper-subgroup residue; see Theorem 5.3 and Remark 5.6.)

The task is therefore not to determine whether a kernel exists algebraically — it does, as \mathbb{Z}_7 , under the strong model — but whether any part of it survives admissibility and the operational quotient as a nontrivial holonomy class. This reformulates the transport programme as a question about the *physical content* of the kernel rather than its existence. The locality, transport-holonomy, and operational-quotient analyses are, in this light, coordinated attempts to determine which branch of the fork is realized — equivalently, to compute where in the chain $N_{\text{obs}} \subseteq N_{\text{hol}} \subseteq N_{\text{alg}}$ the kernel is actually cut off.