Transport Holonomy and the σ-Component of Interface Dynamics

The Flat-Transport Class, the Closed-ness of the Matching Cochain, and What the Interface-Orientation Identification Actually Decides

Keith Taylor VERSF Theoretical Physics Programme versf-eos.com

---

General Reader Summary

The Gate-2 verdict passed exactly when a certain orientation twist of the vacuum is absent — written [s]_vac = 0 — and that condition stands on its own, needing nothing from the rest of the machinery. The verdict also noted that another rule already in the framework, *loop-consistency* (any loop you can shrink to a point closes up trivially), *might* save labour in checking that condition, by disposing of all the shrinkable loops for free. Whether it does was deposited as an open assumption: it depends on whether the orientation information loop-consistency constrains (call it σ) is the same object as the twist the verdict cares about (call it s).

This paper settles the *structure* of that question, even though it cannot settle the question itself without a calculation from inside the interface construction. The structural findings are sharper than the earlier framing, and one of them overturns it.

First, there is a small lemma that the earlier treatment used without stating, and stating it changes everything downstream. Loop-consistency, applied through the orientation read-out, says exactly one thing in precise terms: σ is *closed* — a technical condition meaning it has no local twist. Once that is on the table, the labour-saving shortcut turns out to have nothing to do with σ versus s at all. The shortcut works exactly when **s itself is closed**, full stop — and that is the ordinary state of an orientation twist, guaranteed by the geometry of how the substrate is built whenever the matching twist is the standard one. In that ordinary case the shortcut is free, and the whole σ-versus-s question is idle for it. The identification matters only in the special case where s's closed-ness is *not* guaranteed by the geometry — and there, its only job is to *prove* s closed, after which a completely standard argument does the actual work. So the two "routes" the earlier paper drew as alternatives are not alternatives: one of them merely establishes the other's hypothesis.

Second, the earlier paper's neat three-way classification of how σ and s can relate was, in the very case where it was needed, not exhaustive — it quietly assumed the answer to the open question. The honest picture has a fourth case, and the fourth case is precisely "s is not closed," which is the only world where the shortcut genuinely cannot be rescued. We correct this and state the comparison in language that does not presuppose what is open.

Third, there is a genuinely separate question the earlier framing ran together with the first: whether the interface transport carries anything *beyond* orientation. If it does, that surplus is a real new obstruction for Gate 3 — something the entire orientation programme never saw — and it is the natural home of the interesting failure. The clean way to hold all of this is a single object: the *flat transport class*, the holonomy of the full update around vacuum loops. Orientation is its shadow; the surplus is the part of it the shadow misses; and the two questions become "is the shadow equal to s?" and "does the shadow determine the whole?" We adopt that object as the frame and say exactly what must be computed to pin it down — without pretending to have computed it.

The honest summary: the pass criterion is untouched; the shortcut is free whenever s is closed, which is the ordinary case and has nothing to do with the identification; where closed-ness is open, the identification's only role is to supply it, and the target is a cochain relation, not an equality of classes that would presuppose the conclusion; and whether Gate 3 gains a new layer is a separate splitting question about the same transport object. We reduce all of it to one plaquette-level calculation and state what each outcome means.

Abstract

The Gate-2 verdict reduced reachability to $[s]_{vac} = 0$ and isolated a *verification shortcut* — that interface loop-consistency would discharge the bounding-loop content of that check automatically — as conditional on a deposited **interface-orientation identification** (that the orientation component σ of the interface update coincides with the Gate-2 matching cochain s). This paper formalizes the question, corrects a circularity and a non-exhaustive classification in its prior framing, and recasts it on a single object.

The pass criterion is untouched. B5 ⇔ $[s]_{vac} = 0$ under C1–C3, with no input from the interface dynamics. Everything below concerns only the *shortcut* and the *first higher transport-holonomy layer* of Gate 3.

The bridging lemma (Lemma 3.1, [Proven]). With π_{or} a homomorphism (required for σ_γ to be a holonomy), loop-consistency $U_{\partial P} = \mathbb{1}$ on elementary plaquettes gives $\langle \sigma, \partial P \rangle = 0$, hence $\langle \delta\sigma, P \rangle = 0$ on a 2-chain basis, hence $\boldsymbol{\delta\sigma = 0}$: loop-consistency *implies* σ is a cocycle. (The converse is false when there is surplus — $\delta\sigma = 0$ gives only $U_{\partial P} \in \ker \pi_{or}$, not $U_{\partial P} = \mathbb{1}$; the full equivalence is the conjunction $\delta\sigma = 0$ **and** $\kappa_{\partial P} = \mathbb{1}$, Remark 3.1a + Corollary 3.3. Only the forward implication is used downstream.) This was used silently from the prior framing's §5 onward and is stated here as the bridge. It upgrades the discharge from null-*homotopic* to null-*homologous* (bounding) loops — matching the homology content of $[s]_{vac}$ exactly — and it forces the same closed-ness on the residual κ.

The shortcut is "s closed," and the two routes are not parallel (Theorem 4.1, §4–§5, [Proven]). The bounding-loop discharge the shortcut wants is $\langle s, \partial F\rangle = \langle \delta s, F\rangle = 0$ for all F, i.e. **$\delta s = 0$**. So *shortcut* $\Leftrightarrow$ *s closed*, independent of σ. Closed-ness has two possible sources: **geometric** (s is the closed orientation cochain of the substrate — the ordinary case; then the shortcut is free and the identification is idle), or **interface** (σ closed by Lemma 3.1, plus $s - \sigma$ closed, gives $\delta s = 0$). The interface route is therefore not an alternative to the geometric route — its sole function is to establish the geometric route's hypothesis $\delta s = 0$, after which Theorem 4.1 does the discharging.

The honest four-way split (Theorem 5.1, [Proven]). Given σ closed, set $d = s + \sigma$; then $\delta d = \delta s$. The prior "trichotomy" ($d = 0$ / $[d] = 0$ / $[d] \neq 0$) presupposes [d] defined, i.e. $\delta s = 0$ — the very thing open in the residual case. The exhaustive statement adds the fourth case $\delta d = \delta s \neq 0$ (**s not closed**), the only regime where the interface route cannot rescue the shortcut. Within the residual case, hypotheses are stated at **cochain level** — "$s - \sigma$ a cocycle" ($\Longrightarrow \delta s = 0$, shortcut) or the stronger "$s - \sigma$ a coboundary" ($\Longrightarrow [s] = [\sigma]$ once closed, so σ additionally *computes* [s]_vac on non-bounding loops) — never as "[σ] = [s]," which is circular where [s] is not yet defined.

One object: the flat transport class (§6, [Methodological]). Loop-consistency is flatness of U on bounding loops, so the holonomy layer is a flat-transport class $[U] \in H^1(\Gamma\_vac; G)$ valued in the full update group G, with $1 \to N \to G \to \mathbb{Z}_2 \to 1$ (N = ker π_or, the closure-phase channel; $\mathbb{Z}_2$ = orientation). **Axis A** is whether the π_or-image of [U] equals [s]; **Axis B** is whether that image determines [U] (whether the extension splits on the relevant loops). The axes are independent *once σ is a well-posed cochain* — but that well-posedness (π_or multiplicative across composition) is itself an Axis-B decoupling condition (no ω-feedback into orientation), so the body's prior claim of flat independence is qualified. The surplus class [κ] is legitimate ($\in H^1(\Gamma\_vac; N)$) when N is abelian; nonabelian N puts it in a pointed set with conjugacy caveats — so [κ] is [Conditional on N's structure], not [Proven]. The interesting cell — image correct, extension nonsplit — is where the identification "succeeds" yet Gate 3 gains a real new obstruction; it falls out of the extension picture rather than being assembled by hand.

What remains, as a calculation (§7) with a determinate negative branch (§8). A single plaquette-level evaluation of U's action on the boundary mode — assuming update **homogeneity** across plaquettes, else the full cochain is needed — extracts the cochain relation of $s - \sigma$ and the residual κ, placing the construction in the §7 grid. We do not perform it: it requires the interface construction's explicit boundary-mode action. If that action is not defined to the needed precision (items N1–N3), the determinate finding is that the construction is underspecified at exactly those points, and N1–N3 are simultaneously what pin down the group G of §6.

Labels: [Proven], [Conditional], [Methodological], [Open computation], [Open]. The identification is reduced, not discharged.

Contents

---

1. Position Relative to the Gate-2 Verdict

The Gate-2 verdict returned, unconditionally under C1–C3,

$B5 \iff [s]_{vac} = 0 \iff$ Gate 2 PASS on the vacuum,

with s the $\mathbb{Z}_2$ matching cochain on shared faces of the substrate 2-complex, $[s]_{vac} \in H^1(\Gamma_{vac}; \mathbb{Z}_2)$ its restriction to the vacuum sector Γ_{vac}, and no interface input in the criterion. We do not modify this. Everything here sits strictly downstream and conditions nothing in the verdict.

The verdict additionally isolated a *verification shortcut* — loop-consistency ($U_{\partial\gamma} = \mathbb{1}$ on contractible loops) discharging the bounding-loop content of $[s]_{vac} = 0$ automatically — flagged as conditional on a deposited **interface-orientation identification**. This paper concerns only that shortcut and the holonomy layer of Gate 3 it touches. We make the corrections the prior framing requires: a bridging lemma it used unstated (§3), the relocation of the shortcut to "s closed" (§4), an honest exhaustive comparison stated without circularity (§5), and a single organizing object that unifies the two questions and exposes where the interesting failure lives (§6).

---

2. The Two Structures, and the Orientation Read-out as a Homomorphism

2.1 The Gate-2 matching cochain

Gate 2 assigns each shared face (H, H′) a sign $s_{HH'} \in \{\pm 1\}$ (Theorem 6.1 of the verdict paper). The loop holonomy of a cycle γ is $\mathcal{H}_s(\gamma) = \prod_{(HH') \in \gamma} s_{HH'}$, and $[s] \in H^1(\Gamma_{hub}; \mathbb{Z}_2)$ is its class — *when s is closed*. We write s additively as a $\mathbb{Z}_2$ 1-cochain ($+1 \leftrightarrow 0$, $-1 \leftrightarrow 1$), so

"$\mathcal{H}$_s(γ) = +1" reads "⟨s, γ⟩ = 0," and the closed-ness of s is the condition δs = 0. Whether s is closed is not assumed in this paper; it is the pivot (§4).

2.2 The interface transport and the orientation read-out

The interface framework assigns each elementary update an operator U_{ij} on a state (σ, ω) — σ orientation, ω closure parity — composing around a loop γ to a transport operator U_γ valued in the full update group G. The **orientation read-out** is a map π_or : G → $\mathbb{Z}_2$.

Requirement 2.1 (π_or is a homomorphism) [Open interface fact, pending §8/N2 — not standing background]. For the *orientation holonomy* σ_γ := π_or(U_γ) to be a well-defined $\mathbb{Z}_2$-valued holonomy at all — composing correctly around concatenated loops — π_or must satisfy π_or(U U′) = π_or(U) π_or(U′). It is the minimal condition under which "the orientation component of transport" is a cochain rather than a label. We assume it throughout but flag it as a *substantive open interface fact*, not safe framework background: whether the interface read-out is multiplicative is part of the open item N2 (§8), and Lemma 3.1, the §4 interface route, and all of §5 are conditional on it. Its failure is not benign — it is exactly ω-feedback into the orientation channel, so that the *existence* of a well-posed σ-cochain is itself a decoupling condition of Axis-B type (§6) — and where it fails, σ is not a cochain at all and the comparisons of §4–§5 are not posable until one works with the full transport object [U] over G directly (Remark 6.2).

We write N := ker π_or for the closure-phase channel and κ_γ for the N-component of U_γ, so U_γ is determined by the pair (σ_γ, κ_γ) via the extension 1 → N → G → $\mathbb{Z}_2$ → 1.

---

3. The Bridging Lemma: Loop-Consistency ⟺ σ Closed

The prior framing treated [σ] as a well-defined element of H^1 from the point of comparison onward, without establishing that σ is a cocycle. That step is a lemma, and stating it is what makes §4–§6 correct.

Lemma 3.1 (loop-consistency implies σ is a cocycle) [Proven, modulo Requirement 2.1 — open, see §8/N2]. Let π_or be a homomorphism (Requirement 2.1) and let the elementary plaquettes {P} — the shared faces, single 2-cells — be a basis of 2-chains C_2, with ∂P the boundary 1-cycle of P. Then interface loop-consistency, U_{∂P} = 𝟙 on every plaquette, **implies** δσ = 0.

Proof. π_or(U_{∂P}) = π_or(𝟙) = +1, so ⟨σ, ∂P⟩ = 0 for every plaquette face P. By the boundary–coboundary adjunction, ⟨σ, ∂P⟩ = ⟨δσ, P⟩, so ⟨δσ, P⟩ = 0 on a basis of C_2, whence δσ = 0. ■

Remark 3.1a (the converse fails when there is surplus; the full equivalence is a conjunction) [Methodological]. The implication is one-directional, and deliberately so. The converse — δσ =

$0 \Longrightarrow U_{\partial P} = \mathbb{1}$ — is *false* whenever $N = \ker \pi_{or}$ is non-trivial: $\delta\sigma = 0$ gives only $\langle \sigma, \partial P \rangle = 0$, i.e. $\pi_{or}(U_{\partial P}) = +1$, i.e. $U_{\partial P} \in N$ — the operator lies in the closure channel, not that it is the identity. The orientation *channel* acts trivially; the *operator* need not. Recovering full loop-consistency requires the κ-half as well:

loop-consistency ($U_{\partial P} = \mathbb{1}$) ⟺ ($\delta\sigma = 0$ **and** $\kappa_{\partial P} = \mathbb{1}$),

the σ-conjunct being this lemma and the κ-conjunct Corollary 3.3. We use only the forward implication "loop-consistency ⟹ $\delta\sigma = 0$" downstream, so nothing rests on the false converse; but the σ-half alone is *not* equivalent to full loop-consistency, and is not labelled as such. The gap between the two halves is exactly the surplus the rest of the paper exists to discuss.

Corollary 3.2 (the discharge reaches all bounding loops, via the lemma) [Proven]. With $\delta\sigma = 0$, $\langle \sigma, \gamma \rangle = \langle \sigma, \partial F \rangle = \langle \delta\sigma, F \rangle = 0$ for *every* bounding (null-homologous) loop $\gamma = \partial F$ — not merely the contractible (null-homotopic) ones the axiom mentions. Since the residue $[s]_{vac}$ pairs s with $H_1(\Gamma_{vac}; \mathbb{Z}_2)$, its bounding content is exactly the null-homology that Lemma 3.1 covers. The match between the axiom's reach and the residue's content is therefore a consequence of σ being genuinely closed — **not** of "contractible" and "bounding" being interchangeable (null-homotopic ⟹ null-homologous, not conversely). We use "bounding / null-homologous" hereafter where the prior framing said "contractible."

Corollary 3.3 (κ is trivial on plaquettes — the κ-conjunct) [Proven]. Full loop-consistency $U_{\partial P} = \mathbb{1}$ forces, in addition to $\delta\sigma = 0$ (Lemma 3.1), the N-component $\kappa_{\partial P} = \mathbb{1}$ — this is the κ-half of the equivalence in Remark 3.1a. Hence $\gamma \mapsto \kappa_\gamma$ is trivial on plaquette boundaries and extends to a closed N-valued holonomy on bounding loops, which is what makes the surplus class of §6 well-defined (when N admits a class — §6.3).

---

4. The Shortcut Is Exactly "s Closed"

Theorem 4.1 (the shortcut is the closed-ness of s, and nothing else) [Proven]. The bounding-loop content of the check $[s]_{vac} = 0$ is discharged automatically — "bounding loops free" — **iff** $\delta s = 0$.

Proof. "Bounding loops free" is $\langle s, \gamma \rangle = 0$ for every bounding $\gamma = \partial F$, i.e. $\langle s, \partial F \rangle = \langle \delta s, F \rangle = 0$ for all 2-chains F, which holds iff $\delta s = 0$. ■

So the shortcut has no content beyond s being a cocycle; once $\delta s = 0$, only the non-bounding (non-null-homologous) vacuum loops carry $[s]_{vac}$, and Theorem 4.1 of the *prior* draft (now subsumed) does the discharging. The question is therefore solely: **what makes s closed?**

Two sources of closed-ness, not two parallel routes.

- **Geometric route [the ordinary case].** If the substrate's matching cochain is the orientation (w_1) cochain of a standard oriented 2-cell complex — adjacent 2-cells along shared 1-faces with the usual incidence, the case the verdict paper's "oriented 2-cell complex" language describes — then s is a cocycle by the standard characteristic-class construction (w_1 is a well-defined element of H^1). Here $\delta s = 0$ holds by geometry, the shortcut is free, and the interface-orientation identification is **idle for the shortcut**. The σ–s question simply does not arise for it.
- **Interface route [the residual case].** If the matching cochain is *operational* — defined by transport-matching — and its closed-ness is not geometrically guaranteed, then $\delta s = 0$ is open. Here the interface contributes: σ is closed (Lemma 3.1), and if the difference $s - \sigma$ is itself closed (a cocycle), then $\delta s = \delta\sigma + \delta(s - \sigma) = 0 + 0 = 0$, so s is closed. After that, Theorem 4.1 discharges the bounding loops.

Theorem 4.2 (the routes are nested, not parallel) [Proven]. The interface route is a means of establishing the geometric route's hypothesis. Its entire contribution to the shortcut is the conclusion $\delta s = 0$ (obtained from σ closed + $s - \sigma$ closed); the discharge itself is Theorem 4.1, identical in both cases. Hence the shortcut holds iff s is closed, and the identification matters only as *one possible argument* that s is closed — relevant exactly when the geometric route is unavailable.

Proof. Immediate: both routes terminate in $\delta s = 0$, and Theorem 4.1 takes $\delta s = 0$ to the discharge with no further input. The interface route adds the closed-ness derivation; it does not add a different discharge. ■

This is the first correction: the prior §3 drew geometric and interface routes as alternatives. They are not. The shortcut *is* "s closed," and the interface route is one way to prove that, idle whenever geometry already supplies it.

---

5. The σ–s Relation: the Honest Four-Way Split, at Cochain Level

We treat the residual case (s's closed-ness open) and compare the cochains, taking care that the classification is exhaustive over the regime it is meant to cover and that hypotheses are not stated in class language that presupposes the conclusion.

Definition 5.0 (operational s: the matching read-out, distinct from the holonomy read-out) [Methodological]. The residual case is non-vacuous only if "operational s" is a genuinely different cochain from $\sigma = \pi\_or(U)$ while both are interface-derived — otherwise either s is geometric and σ buys nothing (checking $\langle s - \sigma, \partial P\rangle$ collapses to checking $\langle \delta s, P\rangle$ directly, σ's automatic closed-ness uninformative), or s *is* σ by fiat (case A1, nothing to check). The construction supplies exactly this distinct object. The interface dynamics admits **two read-outs of the same update**:

the matching read-out s_{HH′} — the boundary identification the dynamics *enacts* when it glues hub H to H′ across their shared face (which way the transport sews the two boundary frames together); and **the holonomy read-out** $\sigma = \pi\_or(U)$ — the orientation *projection of the accumulated update holonomy* around the plaquette containing that face.

These are *a priori* different cochains: the sign of the sewing the dynamics performs at a face need not equal the orientation the holonomy reads around a loop through it, because the holonomy folds in the full update (including the closure channel) while the matching sign is the local gluing datum alone. So $s \neq \sigma$ as cochains is the generic situation, not A1-by-fiat — the residual case is non-vacuous. Yet both are read off the *same* dynamics, so their difference $s - \sigma$ is not free: it is constrained by how the interface update relates its gluing action to its holonomy on the boundary mode — which is exactly the (N1–N2) content of §8. This is what makes σ's automatic closed-ness *informative* for an operational s: if the interface relation pins $s - \sigma$ to a cocycle (H-weak), then $\delta s = \delta\sigma = 0$ gives s closed with **no** independent geometric closed-ness argument — and an operational s, not being the geometric orientation cochain, has no such argument available to it. The live regime is therefore precise: an interface-derived matching cochain s, distinct from the holonomy projection σ as a cochain, whose difference $s - \sigma$ is governed by the same (N1–N2) data that fix the group G. We pin this here so that §5 is neither vacuous ($s \neq \sigma$ genuinely) nor idle (σ's closed-ness genuinely bears on s, via the constrained difference).

Theorem 5.1 (exhaustive split, given σ closed) [Proven]. Let σ be closed (Lemma 3.1), and let s and σ be cochains on the common shared-face index fixed by N1 (Definition 5.0) — so that their sum is defined. Set $d := s + \sigma \pmod 2$, so $\delta d = \delta s$. Exactly one of the following holds:

(0) s closed ($\delta d = 0$). Then $[d] \in H^1$ is defined, and exactly one sub-case holds: **(A1)** $d = 0$: $s = \sigma$ as cochains. **(A2)** $d \neq 0$, $[d] = 0$: $s - \sigma = \delta f$ is a coboundary; s is closed and $[s] = [\sigma]$. **(A3)** $[d] \neq 0$: s is closed but $[s] \neq [\sigma]$. **(✗) s not closed ($\delta d = \delta s \neq 0$).** Then d is not a cocycle, [d] is undefined, and **[s] is undefined**.

Proof. $\delta d = \delta s$ since $\delta\sigma = 0$. The dichotomy $\delta d = 0$ versus $\delta d \neq 0$ is exhaustive; within $\delta d = 0$, the cochain d is either zero, nonzero-exact, or non-exact, exhausting (A1)–(A3) by the definition of the cohomology class. The classes in (A1)–(A3) are defined precisely because $\delta d = 0$; in case (✗) no class exists. ■

Remark 5.2 (why the prior trichotomy was not exhaustive) [Methodological]. The prior framing's trichotomy was (A1)/(A2)/(A3) — i.e. case (0) only — stated for "arbitrary $\mathbb{Z}_2$ cochains." But [d] requires d closed, and with σ closed that is $\delta s = 0$, i.e. s closed. The trichotomy therefore silently presupposed s closed, which is exactly what is open in the residual case it was introduced to handle. The fourth case (✗) is the omitted regime, and it is the only one where the shortcut's interface route fails (s not closed ⟹ bounding loops not free). The prior framing's intuition — the dangerous world is non-closed s — was right; the class-level trichotomy could not express it, because class language presupposes closed-ness.

Theorem 5.3 (cochain-level hypotheses, and what each buys) [Proven]. In the residual case the legitimate hypotheses are cochain relations, not class equalities:

(H-weak) s − σ is a cocycle. Then $\delta s = \delta\sigma = 0$, so **s is closed** and the shortcut holds (Theorem 4.1). This is *all the shortcut needs*. It does not by itself give $\mathcal{H}_s = \mathcal{H}_\sigma$ on non-bounding loops.
(H-strong) s − σ is a coboundary. Then additionally $[s] = [\sigma]$ (once both closed), so $\mathcal{H}_s(\gamma) = \mathcal{H}_\sigma(\gamma)$ on **every** loop, and σ *computes* [s]_vac in full — non-bounding loops included.

Proof. (H-weak): $\delta(s - \sigma) = 0$ and $\delta\sigma = 0$ give $\delta s = 0$. (H-strong): $s - \sigma = \delta f$ gives, for any cycle γ, $\langle s - \sigma, \gamma\rangle = \langle \delta f, \gamma\rangle = \langle f, \partial\gamma\rangle = 0$, so $\mathcal{H}_s = \mathcal{H}_\sigma$ everywhere. ■

Corollary 5.4 (the corrected reading of the old "A3 ∧ non-closed s") [Proven]. The prior §5/§7 wrote of "A3 conjoined with non-closed s." This is ill-formed: A3 = "$[s] \neq [\sigma]$" presupposes [s] defined, i.e. s closed, so it cannot coexist with s non-closed. The regime actually intended is case (✗): s not closed (s − σ not a cocycle). Moreover A3 is **not** shortcut-fatal — in A3, s is closed, so Theorem 4.1 discharges the bounding loops regardless; A3 only means σ cannot *compute* the non-bounding holonomy (H-strong fails while H-weak's conclusion still holds via s's own closed-ness). The sole shortcut-fatal regime is (✗).

So Axis A's contribution to the shortcut is captured entirely by H-weak (s − σ a cocycle ⟹ s closed), a cochain condition strictly weaker than the prior "σ = s" and weaker even than $[s] = [\sigma]$; the stronger H-strong is needed only for the separate purpose of computing [s]_vac through σ.

---

6. One Object: the Flat Transport Class and Its Two Questions

The two comparisons of §4–§5 (orientation projection vs s) and the surplus question are most cleanly held as image and splitting data of a single object.

6.1 The flat-transport class

By Lemma 3.1 and Corollary 3.3, full loop-consistency is *flatness* of U on bounding loops: $U_{\partial P} = \mathbb{1}$ (both conjuncts of Remark 3.1a). The Gate-3 holonomy layer is therefore the **flat-transport class** of the full update around vacuum loops, valued in the update group G, with the orientation extension

$1 \longrightarrow N \longrightarrow G \xrightarrow{\pi\_or} \mathbb{Z}_2 \longrightarrow 1$, $N = \ker \pi\_or$ (closure-phase channel).

Remark 6.0a (what kind of object [U] is) [Methodological]. When G is abelian the class is a genuine group element, $[U] \in H^1(\Gamma\_vac; G)$. When G is nonabelian — which it is whenever N is, and can be even for abelian N — the same caveat owed to [κ] in Remark 6.3 is owed here: the flat-transport class is a *pointed set*, $\mathrm{Hom}(\pi_1(\Gamma\_vac), G)$ modulo conjugacy after killing the plaquette boundaries, not a group, and "trivial" means conjugate to the basepoint. There is also a level shift to keep honest: the orientation projection σ and the matching cochain s pair against

$H_1(\Gamma\_vac; \mathbb{Z}_2)$ — abelian, an H^1 object — whereas a nonabelian U-holonomy is genuinely a π_1 object. The image map $\pi\_or \circ [U]$ is precisely the abelianization-and-projection that lands the π_1 object in H^1; the comparison with [s] (Axis A) takes place after that descent, while the surplus (Axis B) is the part that does not descend. We write "$[U] \in H^1(\Gamma\_vac; G)$" for brevity and read it as this pointed set in the nonabelian case.

Orientation holonomy $\sigma = \pi\_or \circ [U]$ is the *image* of [U] under the extension (after the descent of Remark 6.0a); the residual κ is its content in N.

6.2 The two questions as image and splitting

Axis A (image). Does the π_or-image of [U] equal [s]? In the residual case this is asked at cochain level (Theorem 5.3): is $s - \pi\_or(U)$ a cocycle ($\Longrightarrow$ shortcut) or a coboundary ($\Longrightarrow$ σ computes $[s]\_vac$)? **Axis B (splitting).** Does the image determine [U]? Equivalently, does the extension split on the relevant loops — is $\kappa\_\gamma = \mathbb{1}$ for all γ (faithful / split), or $\kappa\_\gamma \neq \mathbb{1}$ for some γ (surplus / nonsplit)?

Theorem 6.1 (where the interesting failure lives) [Methodological]. In the (image, splitting) grid, the consequential cell is **image-correct, extension-nonsplit**: $\pi\_or[U]$ matches [s] (the identification "succeeds," shortcut and/or full computation available) yet $\kappa \not\equiv \mathbb{1}$, so [U] carries a surplus the orientation class never sees. That surplus is a genuine first higher transport-holonomy layer of Gate 3 — invisible to Gate 2 and to any framing in which "$\sigma = s$ settles the holonomy layer." It is not assembled by hand; it is the non-split locus of the extension, read off [U] directly.

Proof sketch. $\pi\_or[U] = [s]$ fixes the $\mathbb{Z}_2$-quotient of [U] to the Gate-2 class; non-split means [U] is not the image of any section $\mathbb{Z}_2 \to G$ on those loops, i.e. κ carries holonomy in N that no lift of the image accounts for. The two data are the quotient (image) and the lifting datum: the image fixes the $\mathbb{Z}_2$-quotient, and the surplus is the remaining lifting content — genuine content **not determined by the image alone**, governed by a coset in $H^1(\Gamma\_vac; N)$ (with $H^2(\Gamma\_vac; N)$ controlling existence of a lift). It sits *over* the image, not beside it; this is consistent with Remark 6.2's qualification that the two are separable invariants only once σ is a well-posed cochain. ■

6.3 Two honesty points the object makes unavoidable

Remark 6.2 (the axes are independent only once σ is a cochain — a decoupling that is itself Axis-B-type) [Methodological]. σ is a well-defined cochain iff π_or is multiplicative across composition (Requirement 2.1). That can fail if the closure-parity channel ω feeds back into the orientation channel — and ω-feedback is precisely ker-π_or / surplus content, i.e. Axis-B territory. So whether **Axis A is even posable** depends on an Axis-B-type decoupling (no $\omega \to \sigma$ feedback). The correct statement is therefore: the two comparisons are independent *once σ is a well-posed cochain*, and the existence of that cochain is itself a decoupling (splitting-of-channels) condition. The flat-transport object absorbs this cleanly — when π_or is not multiplicative, one simply works with $[U] \in H^1(\Gamma\_vac; G)$ directly and does not factor out σ at all.

Remark 6.3 (the surplus is a class only when N admits one) [Conditional on N's structure]. Corollary 3.3 makes κ a closed N-valued holonomy, but $[\kappa]$ is a *cohomology class* only when N has the requisite structure: if N is abelian (as a U(1)-like closure-phase channel would suggest), $[\kappa] \in H^1(\Gamma\_vac; N)$ is legitimate and Theorem 6.1 goes through verbatim. If N is nonabelian, the holonomy lives in nonabelian $H^1(\Gamma\_vac; N)$ — a *pointed set*, not a group — and "$[\kappa]$" carries the usual conjugacy caveats (the class is a conjugacy-orbit, "trivial" means conjugate to the basepoint). Either way, **N must be pinned before [κ] is called a class.** Until then Theorem 6.1's surplus is structural, labelled [Conditional on N's structure], not [Proven]. Pinning N is part of the §7 calculation (item N3 / §8).

---

7. What Must Be Computed, and What Each Outcome Means

We do not manufacture the discharge: placing the construction in the §6 grid requires the interface update's action on the boundary mode, which we do not assume. We state the calculation and tabulate consequences.

The calculation [Open computation]. Read from the interface-dynamics construction the action of U on the boundary mode of one elementary plaquette ∂P, and extract:

(C-A) its orientation projection $\pi\_or(U\_\{\partial P\}) \in \{\pm 1\}$, yielding the cochain σ; and compare $s - \sigma$ — is it a coboundary (H-strong: σ computes [s]*vac), a cocycle but not a coboundary (H-weak: s closed, shortcut, but σ does not compute non-bounding holonomy), or not a cocycle (case ✗: s not closed via this route)? **(C-B)** its residual κ{*∂P} ∈ N, the group N itself, and whether $\kappa \equiv \mathbb{1}$ (split / faithful) or $\kappa \not\equiv \mathbb{1}$ (nonsplit / surplus).

Homogeneity assumption [Methodological]. The single-plaquette → loop-level reduction requires the interface update to be **uniform across plaquettes**, so that one plaquette fixes the rule everywhere. If U is site-dependent, (C-A) and (C-B) must be evaluated as full cochains, not at one plaquette. (Note: Lemma 3.1's *closed-ness* conclusion does not need homogeneity — it uses $\langle \sigma, \partial P \rangle = 0$ for every P — but the §7 *determination of the rule from one plaquette* does.)

Table 7.1 (outcome grid). Rows = the $s - \sigma$ relation (Axis A, cochain level); columns = the splitting (Axis B). The shortcut column-independent fact: in rows R1, R2 the shortcut holds (s closed); in row R3 the shortcut fails via the interface route and needs the geometric route or it does not hold.

	Split ($\kappa \equiv \mathbb{1}$)	**Nonsplit ($\kappa \not\equiv \mathbb{1}$)**
R1: $s - \sigma$ coboundary ($[s]=[\sigma]$, s closed)	Shortcut free; σ computes [s]_vac fully; G3-H collapses to orientation — no new layer.	Shortcut free; σ computes orientation fully; **but [κ] is a genuine new Gate-3 layer** (modulo N, Remark 6.3). *The interesting cell.*

	Split ($\kappa \equiv \mathbb{1}$)	Nonsplit ($\kappa \not\equiv \mathbb{1}$)
R2: s − σ cocycle, not coboundary (s closed, [s]≠[σ])	Shortcut free (s closed); σ gives closed-ness only, not non-bounding holonomy; G3-H orientation part read from s directly; no surplus.	Shortcut free; orientation part from s; **and** surplus [κ] is a new layer.
R3: s − σ not a cocycle (s not closed via interface)	Shortcut fails via interface route — bounding loops not free unless s closed geometrically; no surplus.	Shortcut fails via interface route; **and** surplus [κ] — two independent deficits. (κ's closed-ness is axiom-given: Corollary 3.3 derives $\kappa_{\partial P} = \mathbb{1}$ from full loop-consistency, which holds regardless of s, so the surplus is a well-defined closed N-holonomy even when $\delta s \neq 0$ — s's failure does not contaminate it.)

Reading the table. The top-right cell (R1, nonsplit) is Theorem 6.1's interesting case: the identification fully succeeds on orientation, yet Gate 3 acquires [κ]. R3 is the only shortcut-fatal *interface* outcome, and even it is rescued if s is closed geometrically (§4) — so R3 is fatal only in conjunction with the absence of geometric closed-ness. This corrects the prior table, which mislabelled the class-distinct case (here R2) as shortcut-failing; R2's s is closed, so its shortcut holds.

---

8. The Conditional Negative Result: Underspecification of the Interface Action

The §7 calculation presupposes the construction defines U's boundary-mode action precisely enough that $\pi_{or}(U_{\partial P})$, $\kappa_{\partial P}$, and N are definite objects. This may not hold.

The contingency [Open / Methodological]. For §7 to be executable, the interface construction must supply:

(N1) a definite action of U on the **Gate-2 boundary mode** (the spoke amplitude λ_∂), not merely on the abstract (σ, ω) state; **(N2)** an orientation read-out π_{or} realized as a specific homomorphism on that action (Requirement 2.1), so that σ is a cochain comparable to s and not a label — failing which, by Remark 6.2, Axis A is not even posable and one must work with $[U] \in H^1(\Gamma_{vac}; G)$ directly; **(N3)** the structure of the closure channel $N = \ker \pi_{or}$ — in particular whether ω acts trivially on the boundary mode (split, no surplus) or non-trivially (surplus), and whether N is abelian (so [κ] is a class) — which is simultaneously what fixes the group G of §6.

Theorem 8.1 (the negative result is determinate, and it specifies G) [Methodological]. If the construction does not currently supply (N1)–(N3), the identification is not merely undecided but **not yet well-posed**, since σ, κ, N, and G lack referents on the comparison index until (N1)–(N3) fix them. The finding is then exact: *the interface construction is underspecified at precisely*

(N1)–(N3), which are the minimal extensions required before §7 is attemptable — and, in the §6 framing, (N1)–(N3) are exactly the data that make the update group G and its extension $1 \to N \to G \to \mathbb{Z}_2 \to 1$ precise. Discharging the deposit therefore coincides with specifying G.

Proof. §7 needs σ, κ as objects on the boundary-mode index; (N1) supplies the action, (N2) makes σ a cochain (or, failing multiplicativity, sends one to [U] over G), (N3) fixes N and hence G. Absent these, the symbols of §5–§6 have no instances to classify — Theorem 5.1 is valid for any cochains but has nothing to apply to, and [U] has no group to be valued in. The gap is exactly (N1)–(N3) and is identical to the specification of G. ■

We assume neither that (N1)–(N3) hold nor that they fail. Either the construction supplies them and §7 is a finite calculation placing it in Table 7.1, or it does not and §8 names the extension — which is also the construction of G. Both are determinate; neither is a re-deposit.

---

9. Limitations

The pass criterion is untouched; the identification is not discharged. Nothing here modifies B5 ⟺ $[s]\_vac = 0$ (unconditional, under C1–C3). We establish only framework results valid for any cochains, any homomorphic π_or, and any extension G: Lemma 3.1, Theorems 4.1–4.2, 5.1–5.3 with Corollary 5.4, and the structural Theorem 6.1. The which-cell-of-Table-7.1 question is [Open computation], contingent on §8.

Requirement 2.1 (π_or homomorphic) is load-bearing for Axis A's posability. If the orientation read-out is not multiplicative — ω-feedback into orientation — σ is not a cochain and §4–§5 are not posable; one then works with $[U] \in H^1(\Gamma\_vac; G)$ directly (Remark 6.2). Whether 2.1 holds is part of (N2).

The geometric-shortcut result depends on s's closed-ness, which we do not settle. Theorem 4.1 makes the discharge free iff $\delta s = 0$; whether the substrate's matching cochain is the closed geometric orientation cochain (shortcut free, identification idle) or an operational cochain of open closed-ness (interface route in play, §5) is a reading of the substrate construction we flag but do not perform. The dichotomy is exact; its resolution is external to this paper.

The surplus is conditional on N and its persistence is not addressed. [κ] is a class only when N is abelian (Remark 6.3); nonabelian N puts it in a pointed set with conjugacy caveats. Theorem 6.1's surplus is therefore [Conditional on N's structure], pinned by (N3). And even where [κ] exists and is non-trivial, whether it is a *stable* obstruction — surviving the refinement dynamics $\mathscr{R}$ — is the same refinement-stability question that is load-bearing and open for the defect programme, and is not touched here. A surplus class is an obstruction layer; it is not yet shown to be a persistent one.

The update group G is carried as open structure. §6 organizes the question around $[U] \in H^1(\Gamma\_vac; G)$ without specifying G; (N1)–(N3) are precisely what make G concrete (Theorem

8.1). The flat-transport frame is therefore a *reduction target*, exact in form, awaiting the interface data to become computable. Where G is left abstract, statements about splitting and surplus are structural ([Methodological]/[Conditional]), not [Proven].

The framework is taken as given. The cohomology used — the boundary–coboundary adjunction, H^1, cocycle/coboundary, extensions and nonabelian H^1 as a pointed set — is standard and applied within the VERSF/interface framework; the interface axiom (loop-consistency) and the Gate-2 verdict are taken as established. "Settled" means within that framework.

Gate 3's other layers are untouched. Refinement compatibility and higher-overlap (Čech) assembly — the other candidate Gate-3 layers — are not addressed; the proper organizing structure for Gate 3 (the obstruction tower with w_1 at its base) is a separate treatment.

---

10. Conclusion

The Gate-2 verdict deposited whether the interface orientation component σ coincides with the matching cochain s, as the condition on its verification shortcut. This paper settles the *structure* of that question and corrects its prior framing in three load-bearing ways.

First, a lemma the prior treatment used unstated: with the orientation read-out a homomorphism, loop-consistency *implies* σ is a cocycle (Lemma 3.1; the converse fails under surplus, so the full equivalence is the conjunction with $\kappa_{\partial P} = \mathbb{1}$, and only the forward implication is used). Stating it dissolves a contractible/bounding slippage — the discharge reaches all null-homologous loops, matching the homology content of $[s]_{vac}$, because σ is genuinely closed, not because the two notions coincide. With the lemma, the shortcut's content reduces to a single statement: it holds iff **s is closed** (Theorem 4.1). The geometric and interface "routes" are then not parallel — the interface route's only job is to prove $\delta s = 0$, after which the standard argument discharges (Theorem 4.2). In the ordinary case, where s is the closed orientation cochain of an oriented complex, the shortcut is free and the identification is idle for it.

Second, the prior three-way classification was, in the residual case it was meant to handle, not exhaustive: it presupposed [s] defined, i.e. s closed — the very open point. The honest split adds the fourth case, s not closed (Theorem 5.1), which is the only shortcut-fatal regime, and states the residual-case hypotheses at cochain level — "s − σ a cocycle" for the shortcut, the stronger "s − σ a coboundary" for full computation of $[s]_{vac}$ through σ (Theorem 5.3) — never as "[σ] = [s]," which is circular where [s] is not yet defined. The old "A3 ∧ non-closed s" was ill-formed; the intended regime is the fourth case (Corollary 5.4), and the class-distinct case A3 is not in fact shortcut-fatal, since its s is closed.

Third, the two questions — orientation-vs-s and surplus-beyond-orientation — are the *image* and *splitting* of one object: the flat-transport class $[U] \in H^1(\Gamma_{vac}; G)$, with $1 \to N \to G \to \mathbb{Z}_2 \to 1$. Axis A asks whether the image equals [s]; Axis B whether the image determines the class. The interesting failure — identification succeeds, Gate 3 still gains a layer — is the image-correct, extension-nonsplit cell, and it falls out of the extension rather than being posited (Theorem 6.1). The axes are independent only once σ is a well-posed cochain, and that well-posedness is itself an Axis-B decoupling (Remark 6.2); the surplus is a class only when N admits one (Remark 6.3).

In one line:

The shortcut is free whenever s is closed — the ordinary, geometric case, where the identification is idle; where closed-ness is open, the identification's sole job is to supply it, and the target is the cochain relation "s − σ a cocycle," not a class equality that would presuppose the conclusion; and whether Gate 3 gains a new layer is a separate splitting question about one flat-transport class, settled — together with the group G it lives in — by one plaquette-level evaluation, which we specify rather than assume.

The pass criterion stands untouched throughout. The next concrete step is the §7 evaluation against the interface construction — or, equivalently, the §8 specification of G via (N1)–(N3), which is both what would discharge the deposit and what would make the flat-transport frame computable.