Contextual Distinguishability and Coboundary Confinement

Why Finite Pairwise Comparison Does Not Force a Global Closure Frame, and What Would

Keith Taylor

VERSF Theoretical Physics Programme

Abstract

The Gate-3 closure programme reduces occupancy of the $\mathbb{Z}_7$ closure sector to a single containment question: whether the admissibility constraint subspace A is contained in the coboundary subspace B^1. If $A \subseteq B^1$, every admissible offset is a pure relabelling and the closure class vanishes, $\kappa = 0$. If $A \not\subset B^1$, a nontrivial loop-supported closure residue becomes possible.

A common route toward Branch A assumes that finite distinguishability forces all pairwise closure comparisons to reconcile into a single global frame. This paper shows that the implication is not valid under the minimal pairwise reading of finite distinguishability. We define finite pairwise distinguishability (FPD) as local decidability of each comparison relation, and exhibit an explicit $\mathbb{Z}_7$ triangle satisfying FPD while violating coboundary confinement: three locally valid offsets whose loop sum is nonzero, so that no global vertex labelling generates them.

The non-implication is the defensive content. The paper then does positive work the occupancy programme can use directly. First, it locates the honest antecedent: the substrate condition is not bare FPD but FPD together with the energy-conservation balance law (BCB), which forces $\delta\rho = 0$ and so gives $A \subseteq Z^1$ — every admissible offset is *closed*. Closed is not exact; the gap is $Z^1/B^1 = H^1$, exactly where κ lives, so conservation supplies flatness (condition 1 of the closure programme) without forcing coboundary confinement. Second, it identifies the obstruction to $A \subseteq B^1$ as a first cohomological obstruction in the literal sense — the same object $\kappa \in H^1(\Gamma\_vac; \mathbb{Z}_7)$ the topology paper makes room for — and not merely as something analogous to contextuality. Third, and centrally, it proves a *survival criterion*: a closed admissible configuration produces a surviving $\kappa \neq 0$ if and only if it is supported on a cycle that completion leaves open — its loop sum nonzero **and** outside im ∂_2 — and closedness itself forces any nonvanishing-holonomy witness onto such an open cycle. This fuses the admissibility side of Gate 3 to the topology side: occupancy is not "does admissibility permit a non-coboundary configuration?" but "does admissibility permit one on a T1 generator?" Fourth, it scopes the additional principle Branch A requires — a Global Integrability Principle (GIP), the exact-form condition $\rho = d^0a$ — against the existing VERSF axiom set, proving that conservation is *not* GIP

(it gives closed, not exact) and that closure consistency cannot supply it either, leaving Uniform Readout as the sole remaining suspect. The contrapositive then stands as a diagnostic: any future derivation concluding $A \subseteq B^1$ must invoke an exact-form condition strictly stronger than closedness, and the diagnostic locates it.

The result does not prove that Gate 3 survives, and explicitly does not license reading "Branch A is harder to assert" as support for Branch B. It proves that coboundary confinement is not a consequence of finite pairwise distinguishability alone, fixes the exact further condition under which an admissible witness would settle occupancy, and hands that condition forward as a decidable test.

Contents

---

1. Purpose and Epistemic Scope

The Gate-3 chain has reduced the occupancy question to a single containment:

$A \subseteq B^1$?

where

- A is the subspace of admissible offset configurations,
- B^1 is the subspace of offsets that are pure relabellings (coboundaries),
- $A \subseteq B^1$ means every admissible offset is generated by a global frame assignment,
- $A \not\subseteq B^1$ means at least one admissible offset cannot be globally reconciled.

Both A and B^1 are $\mathbb{Z}_7$-linear subspaces of the cochain space $C^1(\Gamma_{vac}; \mathbb{Z}_7)$, as fixed in the closure-connection paper; the containment is therefore subspace containment, not a set-membership accident, and is refuted by exhibiting a single element of the relevant offset space that lies outside B^1.

This paper asks a narrow question and answers it, then extracts the most the answer will bear.

Narrow question: does finite distinguishability alone force $A \subseteq B^1$?

The answer is no. [proven, §5]

Three boundaries on what follows, stated once so the later sections need not repeat them.

This paper does not claim $A \not\subseteq B^1$ in the actual substrate. [conjectural, deferred to occupancy] The countermodel proves a non-implication about FPD; whether the admissibility rules in fact permit a non-coboundary configuration is the occupancy computation, not this paper.

This paper does not support Branch B. Removing a bad argument *for* Branch A makes Branch B *live*, not *supported*. The distinction is load-bearing and is maintained throughout. A reader who comes away believing the closure residue is more likely to exist has over-read the result.

This paper is conditional on topological availability. [conditional on T1] Its content is empty on the T0 branch of the topology paper, where $A \subseteq B^1$ holds automatically; see §10. It speaks only to the admissibility side of the T1 branch.

---

2. Minimal Pairwise Distinguishability

The non-implication is only as meaningful as the definition of its antecedent. A refutation over a strawman reading of finite distinguishability would be worthless. We therefore state the minimal reading explicitly and mark precisely what it does and does not include.

Definition 2.1 (Finite pairwise distinguishability, FPD)

Let Γ be a finite transport graph with vertex set V and oriented edge set E. A *finite pairwise distinguishability structure* on Γ consists of:

1. a finite set of local comparison outcomes, taken here as $\mathbb{Z}_7$;
2. for each oriented edge $u \to v$, a decidable comparison label $\rho_{uv} \in \mathbb{Z}_7$, with $\rho_{vu} = -\rho_{uv}$;
3. a rule by which equality or inequality of comparison labels is decidable on finite data.

FPD is a condition *local to edges*. It asserts that every pairwise comparison is finitely decidable. It does **not** assert the existence of vertex labels $a_u \in \mathbb{Z}_7$ such that every edge label is a difference $a_v - a_u$. Quantification in Definition 2.1 ranges over pairs (u, v); no clause ranges over loops, paths, or global sections.

Remark 2.2 (The strengthening trap, and the two grades of strengthening)

If a stronger version of finite distinguishability is adopted, the conclusion of this paper may change — but everything depends on *which* strengthening, because there are two grades and they sit on opposite sides of the fork.

The first grade is a **closed-form** strengthening: a condition forcing the offset flow to balance around every admissible completed boundary, $\delta\rho = 0$. This is exactly what the energy-conservation result (BCB) supplies — conservation as a local balance law at each fold, coupling composed relations and constraining every contractible loop. It is genuinely "more than pairwise," and the honest antecedent of this paper is therefore not bare FPD but FPD together with BCB-closedness, giving $A \subseteq Z^1$: every admissible offset is a cocycle. This grade does **not** build Branch A into the definition. Closed is not exact (§3, §8): $\delta\rho = 0$ forces the loop sum to vanish on every plaquette boundary, on all of im ∂_2, and says nothing about cycles completion leaves open. The triangle witness of §4 survives this strengthening — it witnesses non-*exactness* within Z^1, not a failure of closedness — provided its cycle is open, which §7 shows closedness itself forces.

The second grade is an **exact-form** strengthening: a condition forcing every admissible offset to descend from a single global vertex labelling, $\rho = d^0 a$, i.e. $A \subseteq B^1$ directly. This *does* build Branch A into the definition; adopting it would make the entailment a tautology rather than a discovery. The Global Integrability Principle of §8 is precisely a strengthening of this second grade, and the burden §8 isolates is whether any admissibility rule supplies it.

The honest content of the present result is therefore this: under FPD strengthened by BCB-closedness — the actual substrate condition — coboundary confinement does not follow, because closedness is not exactness and the gap between them is $Z^1/B^1 = H^1$. Any opponent who recovers confinement must do so by exhibiting an exact-form (second-grade) principle and arguing it is forced. Conservation is not such a principle; that is the proven negative of §8.

---

3. Coboundary Confinement

Let Γ have vertices V and oriented edges E. An *offset assignment* is a map

$\rho : E \to \mathbb{Z}_7$, $\rho_{vu} = -\rho_{uv}$.

The assignment is a *coboundary* if there exists a vertex labelling

$a : V \to \mathbb{Z}_7$

with

$\rho_{uv} = a_v - a_u \pmod 7$ for every oriented edge $u \to v$.

In the cochain language, $\rho \in B^1$ iff $\rho = d^0 a$ for some 0-cochain a, where d^0 is the coboundary operator dual to ∂_1. Coboundary assignments are exactly those generated by one global frame-numbering of all vertices; $B^1 = \operatorname{im} d^0$.

Lemma 3.1 (Loop-sum obstruction). [proven, telescoping]

For any oriented cycle $\alpha \to \beta \to \gamma \to \cdots \to \alpha$, a coboundary assignment has vanishing loop sum. In the triangle case,

$\rho^{\alpha\beta} + \rho^{\beta\gamma} + \rho^{\gamma\alpha} = 0 \pmod 7$.

Proof. Substituting the coboundary form along the oriented triangle $\alpha \to \beta \to \gamma \to \alpha$,

$(a_\beta - a^\alpha) + (a_\gamma - a_\beta) + (a^\alpha - a_\gamma) = 0$,

the labels telescoping to zero identically. The closing edge $\gamma \to \alpha$ contributes $a^\alpha - a_\gamma$, completing the cancellation. Hence any assignment with nonzero loop sum on some cycle is not a coboundary.

Lemma 3.1 is the whole engine. A nonzero loop sum is a witness, computable from edge data alone, that no global labelling generates the assignment.

---

4. The Countermodel — a Finite Distinguishable Triangle That Is Not a Coboundary

Consider the triangle Γ△ with vertices α, β, γ and oriented edges

$\alpha \to \beta$, $\beta \to \gamma$, $\gamma \to \alpha$,

carrying offset labels

$\rho^{\alpha\beta} = 1$, $\rho^{\beta\gamma} = 1$, $\rho^{\gamma\alpha} = 1$, each in $\mathbb{Z}_7$.

Γ△ satisfies FPD. Each edge carries a single element of $\mathbb{Z}_7$; each comparison is finite and decidable; equality of labels is decidable on finite data. Definition 2.1 is satisfied with nothing left over.

Γ△ is not a coboundary. The loop sum is

$\rho^{\alpha\beta} + \rho^{\beta\gamma} + \rho^{\gamma\alpha} = 1 + 1 + 1 = 3 \pmod 7$,

and $3 \neq 0$ in $\mathbb{Z}_7$. By Lemma 3.1 the assignment is not of coboundary form: there are no vertex labels $a^\alpha, a_\beta, a_\gamma \in \mathbb{Z}_7$ with

$\rho^{\alpha\beta} = a_\beta - a^\alpha$, $\rho^{\beta\gamma} = a_\gamma - a_\beta$, $\rho^{\gamma\alpha} = a^\alpha - a_\gamma$,

since the three equations sum to 0 on the left identically while the labels sum to 3 on the right. The witness is the offset 1-cochain itself, an element of $C^1(\Gamma\triangle; \mathbb{Z}_7)$ lying outside $B^1 = \operatorname{im} d^0$.

Γ△ is therefore locally distinguishable yet not globally integrable. It is a model of the antecedent and a counterexample to the consequent.

---

5. The Non-Implication

Proposition 5.1 (FPD does not force coboundary confinement). [proven, by countermodel]

Finite pairwise distinguishability does not imply $A \subseteq B^1$.

Proof. The triangle Γ△ of §4 satisfies FPD (Definition 2.1) and is not a coboundary (loop sum $3 \neq 0$, Lemma 3.1). It is thus a finite pairwise distinguishability structure whose offset cochain is an element of C^1 outside B^1. The existence of a single FPD structure that is not coboundary-confined refutes the universal implication "FPD ⟹ every admissible offset is a coboundary." Hence FPD alone does not entail $A \subseteq B^1$.

The proposition is exactly a non-implication and nothing more. It does not assert that Γ△ is admissible; §7 fixes the further condition under which an admissible analogue would settle occupancy.

---

6. The Cohomological Obstruction Is the Contextuality Obstruction

The structure of §§3–5 is not merely *reminiscent* of quantum contextuality; it is the same mathematical object, and saying so precisely buys more than an analogy buys.

In the sheaf-theoretic account of contextuality, a family of locally consistent measurement outcomes fails to extend to a global assignment exactly when a certain first cohomological obstruction is nonzero — local sections that do not glue to a global section. The obstruction lives in a first cohomology group of the cover; its nonvanishing is the precise statement that local data are contextual.

The Gate-3 obstruction is the same in kind. Local comparison data are the edge offsets ρ_{uv}; the global section is the vertex labelling a; gluing the local data into a global section is precisely writing $\rho = d^0a$. The obstruction to doing so is the image of ρ in

$H^1(\Gamma\_vac; \mathbb{Z}_7) = Z^1 / B^1$,

and it is nonzero exactly when ρ is a cocycle that is not a coboundary — exactly the loop-sum survival of Lemma 3.1, lifted to homology class. The class is the κ of the closure-connection paper.

Proposition 6.1 (Identification of obstructions). [proven, definitional once the dictionary is fixed]

Under the dictionary

local comparison outcome ↔ edge offset ρ_{uv}, global section ↔ vertex labelling a, gluing ↔ $\rho = d^0a$,

the contextual obstruction to global section and the Gate-3 obstruction to coboundary confinement are the same first cohomology class $\kappa \in H^1(\Gamma\_vac; \mathbb{Z}_7)$. A configuration is contextual in the sheaf sense iff its offset cochain represents a nonzero κ iff it is not coboundary-confined.

This is an identification, not a metaphor, and it has a consequence the metaphor would not deliver: the classification machinery of contextuality transfers. The distinction between *logically* contextual scenarios (some local assignment already fails to extend) and *strongly* contextual ones (no global assignment is consistent with the support at all) becomes a distinction between grades of non-coboundary admissible configuration. Which grade the substrate can realise — if any — is a question the occupancy paper can now ask in borrowed, already-developed vocabulary rather than building the taxonomy from scratch. We flag this as an import worth pursuing [conjectural as to which grade applies]; what is proven here is only the identification of the obstruction, which is what makes the import legitimate.

The correct lesson from quantum theory is therefore specific. It is not a loose appeal to entanglement or gauge redundancy — the former is not a global-frame failure of this kind, the latter is if anything a surfeit of global frames rather than the absence of one. It is the exact fact that local contextual data need not admit a global section, and that the obstruction to their doing so is a first cohomology class.

---

7. The Survival Criterion — Fusing Admissibility to Topology

Proposition 5.1 leaves a gap that must not be papered over: a non-coboundary offset is necessary for occupancy but not sufficient. A triangle with nonzero loop sum settles nothing if completion bounds that triangle — if the cycle carrying the offset is itself filled by a plaquette, its class dies in the topology regardless of admissibility. The occupancy witness must live where the topology leaves room. This section makes that precise and, in doing so, ties the admissibility side of Gate 3 to the topology side as a single condition.

Recall from the topology paper the rank-deficit picture: cycles live in $Z_1 = \ker \partial_1$ of dimension μ, completion bounds the subspace $\mathrm{im}\, \partial_2 \subseteq Z_1$, and the surviving sector is

$H_1(\Gamma\_vac; \mathbb{Z}_7) = Z_1 / \mathrm{im}\, \partial_2$,

dual to the H^1 in which κ lives. A nonzero loop sum certifies that an offset cochain evaluates nontrivially on some cycle; whether that nontriviality *survives* depends on whether the cycle is bounded by completion.

Theorem 7.1 (Survival criterion). [proven, given the topology paper's rank-deficit picture]

Let $\rho \in C^1(\Gamma\_vac; \mathbb{Z}_7)$ be an admissible offset cochain, and suppose ρ is a cocycle ($\rho \in Z^1$). This closedness is not an extra hypothesis on the occupancy side: it is supplied upstream by the energy-conservation balance law (BCB), which forces $\delta\rho = 0$ — no net imbalance around any admissible completed boundary — and so establishes $A \subseteq Z^1$ as a derived input rather than an assumption. Then ρ represents a surviving class $\kappa \neq 0$ in $H^1(\Gamma\_vac; \mathbb{Z}_7)$ if and only if there exists a cycle $z \in Z_1$ with

$\langle \rho, z \rangle \neq 0$ in $\mathbb{Z}_7$ **and** $z \notin \mathrm{im}\, \partial_2$, equivalently, iff the loop sum of ρ is nonzero on some cycle that completion leaves open.

Proof. By the universal-coefficient identification over the field $\mathbb{Z}_7$, $H^1 \cong \mathrm{Hom}(H_1, \mathbb{Z}_7)$, so a cohomology class is nonzero iff it pairs nontrivially with some homology class. The homology classes are represented by cycles $z \in Z_1$ modulo im ∂_2. The pairing $\langle\rho, z\rangle$ is well-defined on H_1 when $\rho \in Z^1$: for a homologous cycle $z' = z + \partial_2 c$, the adjointness of ∂_2 and the coboundary δ gives $\langle\rho, z'\rangle = \langle\rho, z\rangle + \langle\rho, \partial_2 c\rangle = \langle\rho, z\rangle + \langle\delta\rho, c\rangle = \langle\rho, z\rangle$, since ρ is a cocycle ($\delta\rho = 0$). The value therefore depends only on the homology class of z. Hence $\kappa = [\rho]$ is nonzero in H^1 iff $\langle\rho, z\rangle \neq 0$ for some z whose class in H_1 is nonzero, i.e. some $z \in Z_1$ with $z \notin \mathrm{im}\ \partial_2$. The loop sum of ρ around a cycle z is exactly $\langle\rho, z\rangle$; nonzero loop sum on an open cycle is precisely the stated condition.

Corollary 7.2 (The occupancy question, sharpened). [proven]

A non-coboundary admissible configuration settles occupancy positively if and only if it is supported on a generator of $H_1(\Gamma\_vac; \mathbb{Z}_7)$ — a T1 surviving loop. Equivalently, occupancy is not the question

"does admissibility permit a non-coboundary configuration?"

but the strictly stronger question

"does admissibility permit a non-coboundary configuration on a cycle completion leaves open?"

Proof. Immediate from Theorem 7.1: a non-coboundary offset on a cycle inside im ∂_2 has $\langle\rho, z\rangle$ either zero on all open cycles or nonzero only on bounded ones, and represents $\kappa = 0$; only support on an open cycle yields $\kappa \neq 0$.

This is the central upgrade. It removes the daylight between "non-coboundary" and "surviving residue," and it fuses the two Gate-3 papers into one pipeline. The topology paper decides whether open cycles exist at all (T0 versus T1, the existence of generators of H_1). The present criterion decides what an admissible witness must do to occupy one (carry nonzero loop sum on such a generator). Neither paper alone closes Gate 3; together they reduce it to a single conjoined test:

Gate 3 occupied $\Leftrightarrow$ ($\exists$ open cycle: T1) $\wedge$ (admissibility permits nonzero loop sum on it).

The first conjunct is the topology paper's rank computation. The second is the occupancy paper's semantic bridge. Theorem 7.1 is the hinge that makes them one condition rather than two adjacent ones.

Remark 7.3 (Closedness forces open-cycle support — the witness cannot hide)

The conservation law and the survival criterion are the same constraint seen from two sides, and noticing this removes a hypothesis. $\delta\rho = 0$ says the loop sum of ρ vanishes on every plaquette boundary — on all of im ∂_2. A surviving κ needs nonzero loop sum on a cycle *outside* im ∂_2. So once BCB supplies closedness, any offset with nonzero holonomy at all is *forced* onto an open cycle: it cannot carry a nonzero loop sum on a completed boundary, because conservation forbids exactly that. The open-cycle support of Corollary 7.2 is therefore not a side condition the witness must be arranged to satisfy — it is an automatic consequence of closedness together with nonvanishing holonomy. In particular the degenerate reading of the §4 triangle, in which its 3-cycle is itself a completed plaquette, is ruled out by conservation: a completed boundary has loop sum zero by $\delta\rho = 0$, so the loop sum $3 \neq 0$ certifies the triangle's cycle is open. Closedness does the bookkeeping that would otherwise have to be imposed by hand.

Remark 7.4 (Torsion case)

By the torsion mechanism of the topology paper, an open cycle may be a 7-torsion class — a loop closing only after seven traversals — invisible rationally but contributing to $H_1(\mathbb{Z}_7)$. The survival criterion applies unchanged: $\langle\rho, z\rangle$ is computed over $\mathbb{Z}_7$, and $z \notin \operatorname{im} \partial_2$ is tested over $\mathbb{Z}_7$, so a torsion generator is a legitimate support for a surviving κ exactly as a free generator is. The witness triangle of §4 is free (its open cycle, were it open, would be rationally visible); the substrate's surviving generator may instead be torsion, and the criterion covers both without modification.

---

8. Scoping the Global Integrability Principle

Proposition 5.1 shows Branch A requires a principle beyond FPD. Naming the principle is cheap; scoping where it could come from is the work that helps the sequel.

Definition 8.1 (Global Integrability Principle, GIP)

GIP is the assertion that every admissible pairwise comparison arises from a global vertex labelling:

for every admissible ρ, there exists $a : V \to \mathbb{Z}_7$ with $\rho = d^0 a$.

Equivalently, $A \subseteq B^1$. GIP is exactly the content Branch A needs and exactly the content FPD lacks.

The question for the programme is whether GIP is *derivable* from the existing VERSF axiom set or must be *posted* as new. We scope this, candidate by candidate, without settling it — settling it is the occupancy paper's task, but the short list and the reasons are handed forward here. [conditional / conjectural as marked per candidate]

Uniform Readout (UR). UR requires that readout assign outcomes uniformly across admissible configurations. This is the most natural candidate to *look* like GIP, because uniformity of readout resembles existence of a single global frame. But the resemblance is not yet an entailment: UR constrains the readout map's action on configurations, not the existence of a global *integrating* labelling for the offsets. UR could hold while offsets remain contextually frustrated, provided readout is uniform on each context separately. Whether UR forces a *global* section or only *per-context* sections is precisely the logical-versus-strong contextuality distinction of §6, and is the single most important thing to check. [conditional — UR is the prime suspect and the prime place GIP could hide]

The data/dynamics cut. The cut separates committed data from reversible dynamics. It governs which structures are fixed facts versus which are still labels in play. It does not obviously assert that the fixed data are globally integrable; a committed datum can be a contextual comparison as readily as a global coordinate. The cut therefore does not appear to supply GIP. [conjectural — no evident route, but not excluded]

Closure consistency at hubs. Hub completion enforces minimal closure cycles. This is a *local* consistency condition and, by the §6.1 analysis of the topology paper, locality does not force global integrability — dense local closure can still leave noncontractible or torsion cycles unbounded. Closure consistency therefore cannot supply GIP on its own; if it could, the topology paper's T1 branch would be empty, which it is not. [proven not to supply GIP unaided, given the topology paper]

Energy conservation (BCB). This is the candidate most likely to be mistaken for GIP, and the mistake is exactly the closed-versus-exact conflation. Energy conservation, emerging as a balance law on commitment events, forces $\delta\rho = 0$ — no net imbalance around any admissible completed boundary. That is the **closed-form** condition, $\rho \in Z^1$, condition (1) of Proposition 6.1 (flatness). It is necessary for GIP but strictly weaker than it: GIP is the **exact-form** condition $\rho = d^0 a$, and closed does not imply exact when the topology carries an open cycle, the gap being precisely $Z^1/B^1 = H^1$. So energy conservation supplies a *different* one of the three Gate-3 conditions — flatness — and is structurally incapable of supplying GIP: a balance law constrains differences around bounding loops, not the existence of a single global potential. The Aharonov–Bohm situation is the standing illustration — a field closed everywhere accessible, energy conserved, yet the potential not globally single-valued and the holonomy around the inaccessible flux physically real. Conservation gives flatness without killing occupancy. [proven not to supply GIP — it supplies flatness instead, closed not exact]

Internality / decomposition-independence (IA, PC). These govern invariance under choice of decomposition. They constrain how comparisons transform, not whether comparisons globalize. They are orthogonal to GIP as far as can be seen. [conjectural — orthogonal]

The scoping yields a sharp residue. Of the existing axioms, **UR is the only candidate that could plausibly entail GIP**, and whether it does reduces to the logical-versus-strong contextuality question of §6 — does UR force a global section, or only local ones? Closure consistency demonstrably cannot supply GIP unaided. Energy conservation demonstrably does not supply it either — it supplies flatness, a different condition, and the conflation of the two is the closed-versus-exact error. The remaining axioms appear orthogonal. So either GIP descends from UR — in which case Branch A holds and the occupancy paper's task is to prove that descent — or GIP must be posted as a genuinely new principle, in which case its postulation is a visible, isolable assumption rather than a buried one. Either way the burden is now located at a single axiom, and the strongest-looking decoy — conservation — has been removed from the list for a structural reason rather than by elimination-fatigue.

---

9. The Contrapositive as a Programme Diagnostic

Proposition 5.1 read forwards says FPD does not force A. Read as a contrapositive it is a standing check on every future paper in the chain.

Proposition 9.1 (Diagnostic). [proven, contrapositive of 5.1]

Any derivation concluding $A \subseteq B^1$ must invoke, somewhere in its premises, a condition strictly stronger than FPD.

Proof. If a derivation concluded $A \subseteq B^1$ from FPD alone, it would establish the implication refuted by Proposition 5.1. Hence any valid such derivation uses a premise not implied by FPD.

The diagnostic use is operational. When a future paper claims Branch A, one locates the extra assumption by finding the step that exceeds pairwise decidability — the first place the argument quantifies over loops, paths, or global sections rather than over pairs. By §8 the prime location to inspect is wherever Uniform Readout enters. Proposition 9.1 turns the present result from a statement about one fork into a reusable instrument: it guarantees the extra assumption exists and tells the reader what kind of step to look for.

---

10. Relation to Topological Availability

This argument matters only on the topologically available branch, and its conditionality must be stated, not assumed.

If the vacuum transport complex has trivial first cohomology — the T0 branch of the topology paper, rank $\partial_2 = \mu$ — then $H^1(\Gamma\_vac; \mathbb{Z}_7) = 0$, every cocycle is a coboundary, $A \subseteq B^1$ holds automatically, and $\kappa = 0$ regardless of admissibility. On T0 the present paper is vacuous: there is no open cycle for an admissible witness to occupy, and Theorem 7.1's survival condition is unsatisfiable. The fork of §11 collapses.

The present result therefore reads conditionally:

Given T1 — a topology supporting a nontrivial $\mathbb{Z}_7$ class — finite pairwise distinguishability alone does not force admissible offsets into coboundaries, and an admissible non-coboundary configuration occupies the sector iff it is supported on a surviving generator (Theorem 7.1).

The paper does not replace the topology criterion; it supplies the admissibility side of the same gate, joined to the topology side by the survival criterion, with flatness fed in from the conservation result. The pipeline is explicit:

energy conservation (BCB): $\delta\rho = 0$, so $A \subseteq Z^1$ (flatness, condition 1) ↓ topology paper: is there an open cycle? (T1 versus T0, condition 2) ↓ if T1 this paper: can admissibility carry nonzero loop sum on it? (Theorem 7.1, §8, condition 3) ↓ if yes $\kappa \neq 0$, sector occupied.

The three stages are the three conditions of Proposition 6.1, now distributed across three results: conservation discharges flatness, the topology paper decides availability, and the occupancy question is escape from B^1 within the Z^1 that conservation guarantees.

---

11. Reframing the Gate-3 Occupancy Problem

The result reframes the occupancy question and the reframing is now backed by a criterion rather than a slogan.

Old framing:

Does finite distinguishability collapse the offset space into coboundaries?

Corrected framing:

Does admissibility add a global-integrability condition (GIP) beyond finite pairwise distinguishability — and if not, does it permit a non-coboundary configuration on a cycle completion leaves open?

The two branches, with their now-precise conditions.

Branch A — global integrability. Admissibility entails GIP (by §8, plausibly via Uniform Readout; otherwise by new postulate). Then $A \subseteq B^1$ and $\kappa = 0$. The sector is unoccupied by relabelling-triviality.

Branch B — contextual comparison. Admissibility permits offsets that are locally decidable but globally frustrated, and — by Theorem 7.1 — permits at least one such offset on a surviving generator of H_1. Then $A \not\subseteq B^1$ and $\kappa \neq 0$, provided flatness and topology also permit it.

The paper does not decide between the branches. It establishes that Branch A is not forced by FPD, fixes the exact condition (Corollary 7.2) under which a Branch-B witness would count, locates the single axiom (UR) on which Branch A most plausibly turns, and supplies a diagnostic (Proposition 9.1) for auditing any future Branch-A claim. What it does **not** do, and must not be read as doing, is argue that Branch B obtains: the live possibility of contextual frustration is not evidence of it.

Conclusion

The Gate-3 programme has repeatedly encountered the question of whether admissible offsets are confined to coboundaries. This paper proves that confinement does not follow from finite pairwise distinguishability alone, by an explicit countermodel: a $\mathbb{Z}_7$ triangle with offsets 1, 1, 1 has finite decidable comparisons on every edge yet loop sum $3 \neq 0$, so no global vertex labelling generates it. Local decidability does not imply global integrability.

That non-implication is the defensive content. The strengthening rests on one structural distinction — closed is not exact — and runs in four moves. The honest antecedent is FPD together with the energy-conservation balance law (BCB), which forces $\delta\rho = 0$ and so gives $A \subseteq Z^1$: admissible offsets are closed. Closedness is condition (1), flatness, supplied to the occupancy side as a derived input rather than an assumption; it constrains the loop sum to vanish on every completed boundary and so forces any nonvanishing-holonomy witness onto an open cycle, but it does not force exactness, the gap being $Z^1/B^1 = H^1$. The obstruction to coboundary confinement is then identified — not analogized — as the first cohomological obstruction of contextuality, the same class κ the topology paper makes room for, which licenses importing the contextuality taxonomy. The survival criterion of Theorem 7.1 fuses the admissibility side of Gate 3 to the topology side: a closed admissible configuration occupies the sector iff its loop sum is nonzero on a cycle completion leaves open, so occupancy is the conjoined test "open cycle exists (T1) and admissibility carries nonzero holonomy on it." And the Global Integrability Principle that Branch A requires — the exact-form condition $\rho = d^0 a$ — is scoped against the existing axioms with two proven negatives: conservation does not supply it (it gives closed, not exact), and closure consistency does not supply it unaided (locality does not force global section), localizing the burden to a single prime suspect, Uniform Readout. The contrapositive stands as a diagnostic: any future Branch-A derivation must exceed closedness with an exact-form condition, at a step the result tells the reader how to find.

The paper does not prove that the VERSF substrate permits a surviving residue, and does not treat the weakening of Branch A as support for Branch B. It proves that forbidding the residue requires an additional admissibility principle, fixes the exact further condition under which an admissible witness would establish the residue, and hands that condition forward as a decidable test conjoined to the topology criterion. The decisive issue is no longer whether comparisons are finite and decidable. It is whether admissibility requires all such comparisons to globalize — and, if it does not, whether it permits one that fails to globalize on a loop the topology leaves open.

Branch A requires global integrability. Branch B requires contextual comparison on a surviving generator. Finite distinguishability alone decides neither — but the criterion that would decide between them is now explicit.