Uniform Readout and Global Integrability

The Path-Independent Ledger as the Final Input to Gate-3 Occupancy

Keith Taylor

VERSF Theoretical Physics Programme

Summary for the General Reader

This paper finishes locating — though it does not finish answering — the last question in the Gate-3 closure arc. It also corrects a wrong turn taken in a superseded companion paper, which tried to obtain the closure charge by treating it as a "winding" of the continuous quantum phase, reduced to a seven-position value. That route does not match how the framework actually defines the charge, and it is set aside here.

Here is the corrected picture in plain terms. The closure charge — the "twist" the arc has been chasing — is not a leftover from the smooth quantum phase at all. It is a property of how a network of committed comparisons fits together. Picture a web of sites, with each connecting edge carrying a comparison label that says how the two ends differ, on a seven-mark dial. The question is whether all these local comparisons can be reconciled into a single consistent global numbering of the sites — like checking whether a set of "A is 3 ahead of B, B is 2 ahead of C" statements can be made consistent by assigning each person a single position — or whether they are *globally frustrated*, like the famous impossible staircase where every local step looks fine but going all the way around brings you back higher than you started.

A companion paper proved something important: local decidability does not settle this. You can have a web where every single comparison is perfectly definite, yet no global numbering reconciles them — a triangle whose three "+1" steps sum to 3 around the loop instead of 0. So whether the charge survives comes down to whether the framework's rules *force* a global reconciliation or merely permit local consistency. That same companion paper then checked the framework's principles one by one and found that almost none of them force global reconciliation — energy conservation, in particular, only forces a weaker condition. Exactly one principle remained a live candidate: **Uniform Readout**, the rule about how committed outcomes are extracted from the substrate.

This paper takes up that single remaining question and does three things. First, it proves an exact equivalence: Uniform Readout forces the charge to vanish **if and only if** reading out the substrate amounts to consulting a single, path-independent global record — a "ledger" whose value at each site does not depend on the route you took to reach it. Second, it shows this is the very same distinction that separates an impossible-staircase situation from an ordinary one —

what the physics-of-measurement literature calls the difference between merely-local and globally-reconcilable data. Third, it pins down the one concrete property of the readout rule that decides everything: does committing facts ever lay down a *global* frame, or does it only ever commit *local* comparisons between neighbours?

If the framework only ever commits local comparisons — which fits its picture of facts as small, local, irreversible events — then there is no mechanism forcing global reconciliation, and the charge can survive. If instead reading out the substrate consults a single global ledger, the charge vanishes. The paper does not claim to know which; it proves that this is the whole question, and hands forward a clean test — one that becomes decidable as soon as the relevant rule is stated precisely enough. The honesty is the point: rather than manufacture an answer, it shows that the entire arc now rests on one checkable property of one principle.

The technical body makes all of this precise.

Note on This Paper's Relationship to a Superseded Companion Paper

A superseded companion paper, *From Bit Count to Closure Charge*, attempted to derive the closure charge as the mod-7 reduction of an integer winding of the amplitude phase circle U(1), motivated by a connectedness no-go forbidding a continuous map $U(1) \to \mathbb{Z}_7$. That no-go is a correct theorem, but it concerns the wrong object. As fixed in the closure-connection paper and made explicit in *Contextual Distinguishability and Coboundary Confinement*, the closure charge is not a projection of the continuous amplitude phase. It is a cohomology class $\kappa \in H^1(\Gamma\_vac; \mathbb{Z}_7)$ on the vacuum transport complex, with $\mathbb{Z}_7$ entering as the **coefficient group** of the offset cochains — the $K = 7$ constraint count serving as the comparison alphabet — not as a winding period of the amplitude circle. The amplitude U(1) of the measurement layer and the closure $\mathbb{Z}_7$ of the offset complex were never two circles requiring a homomorphic bridge. The present paper supersedes that companion and addresses the actual open input the arc has converged on.

Table of Contents

1. **The Corrected Object** — the closure charge is $\kappa \in H^1(\Gamma\_vac; \mathbb{Z}_7)$ on the vacuum transport *complex*, a contextuality obstruction, not a winding of the amplitude phase.
2. **The Pipeline, and Where It Terminates** — flatness from conservation, availability from topology, the survival criterion, the fixed-complex reading, and the single remaining gate: $A \subseteq B^1$.
3. **Global Integrability and the Section Problem** — restating GIP as the existence of a global section, and the proven negatives that leave Uniform Readout as the sole suspect.
4. **The Equivalence Theorem** — Theorem 4.1: Uniform Readout entails $\kappa = 0$ if and only if readout consults a path-independent ledger (a global section).

Abstract

The Gate-3 occupancy question has been reduced, by *Contextual Distinguishability and Coboundary Confinement*, to a single containment: whether the admissible offset subspace A is contained in the coboundary subspace $B^1 \subseteq C^1(\Gamma\_vac; \mathbb{Z}_7)$. If $A \subseteq B^1$ the closure class vanishes ($\kappa = 0$, sector empty); if $A \not\subseteq B^1$ a surviving residue is possible. That paper proved that finite pairwise distinguishability does not force $A \subseteq B^1$, supplied flatness ($A \subseteq Z^1$) from the energy-conservation balance law, fused admissibility to topology via a survival criterion, and scoped the Global Integrability Principle (GIP: $A \subseteq B^1$) against the VERSF axiom set — obtaining proven negatives for conservation and closure consistency and isolating Uniform Readout (UR) as the sole remaining candidate source of GIP — conditional on that enumeration of principles being exhaustive — with the question reduced to the logical-versus-strong contextuality distinction.

This paper takes up that single residual question. It first corrects a prior misframing: the closure charge is $\kappa \in H^1(\Gamma\_vac; \mathbb{Z}_7)$, a first-cohomology (contextuality) obstruction with $\mathbb{Z}_7$ as cochain coefficient group, not a winding of the amplitude phase circle. It then proves the central equivalence (Theorem 4.1): UR entails $\kappa = 0$ for all admissible configurations if and only if readout consults a path-independent ledger — equivalently, if and only if UR supplies a global section $a : V \to \mathbb{Z}_7$ with $\rho = d^0 a$ for every admissible ρ. This identifies UR $\Longrightarrow$ GIP with ledger path-independence and with non-contextuality in the sheaf sense, characterizing UR-global through three equivalent conditions (GIP, path-independent ledger, non-contextuality) the arc had treated as separate vocabularies — the equivalence of UR-global with the first being essentially definitional, the substance lying in the equivalence of the three. It decomposes UR into a frame component and a comparison component, proving that only the frame component can carry the global force GIP requires, and that the comparison component is satisfied by the contextually frustrated triangle witness. It then identifies the single deciding property: whether commitment ever lays down a global frame (a committed 0-cochain $a : V \to \mathbb{Z}_7$) or only local comparison facts (committed 1-cochain offsets ρ on edges). Both the Branch-A reading (a single path-independent ledger, $\kappa = 0$) and the Branch-B reading (only local committed comparisons, $\kappa \neq 0$ possible) are stated without prejudice. A Local Commitment Bias result (Proposition 6.2) and

its Burden Shift corollary note that, since the primitive ontology is exhaustively local, Branch A must exhibit a global-frame-fixing mechanism the generative chain does not supply — a parsimony asymmetry relative to the established primitives, not a decision, and conditional on the local-ontology premise the test itself checks. The result does not settle occupancy; it proves that occupancy now rests on exactly one well-posed property of the readout axiom — decidable once its quantifier is fixed — and supplies the test.

---

1. The Corrected Object

The closure charge of the Gate-3 sector is the cohomology class

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

where Γ_vac is the vacuum transport complex — a 2-complex whose vertices V are substrate sites, whose oriented edges carry the pairwise comparisons, and whose 2-cells are attached by the closure constraints at hubs (see §2 and §7). Here $C^1(\Gamma\_vac; \mathbb{Z}_7)$ is the space of $\mathbb{Z}_7$-valued offset 1-cochains, $Z^1 = \ker d^1$ is the closed (cocycle) subspace under the coboundary $d^1 : C^1 \to C^2$, and $B^1 = \operatorname{im} d^0$ is the exact (coboundary) subspace under $d^0 : C^0 \to C^1$. The offsets $\rho_{uv} \in \mathbb{Z}_7$ are the decidable pairwise comparison labels on oriented edges, with $\rho_{vu} = -\rho_{uv}$. The 2-cells are essential rather than decorative: on a pure graph (no 2-cells) one would have $C^2 = 0$, $d^1 = 0$, $Z^1 = C^1$, and both the flatness condition of §2 and the torsion of §7 would be vacuous — every offset would be closed automatically and H_1 would be free. It is the closure 2-cells that give flatness content and that carry the $K = 7$ torsion (§7). This complex is fixed in the closure-connection paper and is the operative definition throughout the Gate-3 arc.

Three points fix the object and correct the earlier draft.

First, **$\mathbb{Z}_7$ is the coefficient group of the offset cochains, not a phase period.** The seven values are the comparison alphabet — the $K = 7$ constraint count, repurposed as the set of distinguishable pairwise offsets. The charge does not arise by reducing a continuous amplitude phase modulo anything. The amplitude circle U(1) of the measurement layer (selected by the division-algebra argument, carrying the Born-rule phase) and the closure group $\mathbb{Z}_7$ are different structures playing different roles: the former is the phase of complex amplitudes, the latter the coefficient ring of a cochain complex on the substrate complex. No homomorphism $U(1) \to \mathbb{Z}_7$ is needed or sought; the prior connectedness no-go, while true, concerns an object the framework does not use.

Second, **the charge is a contextuality obstruction in the literal sense.** Under the dictionary {local comparison outcome ↔ edge offset ρ_{uv}; global section ↔ vertex labelling a; gluing ↔ $\rho = d^0 a$}, the obstruction to coboundary confinement is identical to the first-cohomological obstruction to a global section in the sheaf-theoretic account of contextuality (Abramsky–Brandenburger) and its cohomological refinement (Abramsky–Mansfield–Barbosa), from which the logical-versus-strong contextuality terminology used in §3 and §5 is taken. This is an identification, not an analogy: κ is nonzero exactly when the local comparison data fail to glue to a global section.

Third, **occupancy is the containment question $A \subseteq B^1$.** The sector is empty ($\kappa = 0$) iff every admissible offset is a coboundary; it can be occupied ($\kappa \neq 0$) iff some admissible offset is closed but not exact, supported on a surviving cycle.

---

2. The Pipeline, and Where It Terminates

The occupancy question is the conjunction of three conditions, now distributed across three results of the arc and reduced to a single open input.

Condition 1 — Flatness (proven). The energy-conservation balance law (BCB) forces $d^1\rho = 0$ on every admissible offset, so $A \subseteq Z^1$: every admissible offset is closed — its value on the boundary of every 2-cell vanishes. Because Γ_vac carries 2-cells (those attached by ordinary, degree-one closure constraints fill loops and impose real conditions), this is a genuine constraint, not the vacuous condition it would be on a pure graph: conservation supplies flatness as a derived input, not an assumption.

Condition 2 — Availability (topology paper). Whether Γ_vac carries a nontrivial first cohomology — the T1 branch, with open cycles outside im ∂_2 — is the rank-deficit computation of the topology paper. On the T0 branch $H^1 = 0$, every cocycle is a coboundary, and $\kappa = 0$ automatically; the occupancy question is vacuous there. The remainder of this paper is conditional on T1.

Condition 3 — Admissible escape from B^1 (the open input). Given flatness and availability, the sector is occupied iff admissibility permits a closed offset that is not exact, supported on a surviving generator. By the survival criterion, a closed admissible ρ represents $\kappa \neq 0$ iff its loop sum is nonzero on a cycle completion leaves open — and closedness itself forces any nonvanishing-holonomy witness onto such a cycle, so open-cycle support is automatic rather than an extra demand.

The pipeline therefore terminates at one question: does admissibility force $A \subseteq B^1$ (every closed admissible offset is exact), or does it permit a closed-but-not-exact admissible offset? This is the Global Integrability question, and the arc has shown it turns on a single axiom.

The complex is fixed; commitment selects the realized offsets. Sections 1–4 compute cohomology on a fixed Γ_vac, while §6 describes reality as accumulating by local commitment events. These are reconciled by the following reading, which the rest of the paper adopts: Γ_vac is the fixed vacuum transport complex, and commitment does not alter it. Commitment selects which closed offsets are realized — it determines the admissible set $A \subseteq Z^1$ — but it does not change the complex on which Z^1, B^1, im ∂_2, and any torsion are computed. The topology is a property of the fixed vacuum complex; the local commitment events of §6 populate A within that fixed cohomological setting. The well-posedness of the §8 test depends on this: "is $A \subseteq B^1$" is asked on a fixed complex with a fixed B^1. (If instead the complex grew as commitments accrued, H_1 would be dynamical and the framing would require a separate stability-under-refinement argument; the operative reading here is the fixed-complex one, and the §6.1 locality argument concerns which offsets in the fixed C^1 get realized, not growth of the complex.)

---

3. Global Integrability and the Section Problem

Definition 3.1 (Global Integrability Principle, GIP) (definitional). GIP asserts that every admissible offset descends from a single global vertex labelling: for every admissible ρ there exists $a : V \to \mathbb{Z}_7$ with $\rho = d^0 a$, i.e. $\rho_{uv} = a_v - a_u$ for every oriented edge. Equivalently, $A \subseteq B^1$.

GIP is exactly the existence of a **global section** of the comparison data: a single consistent assignment of a $\mathbb{Z}_7$ value to every site such that every edge offset is the difference of its endpoints' values. The triangle witness — three edges forming an **unfilled (open)** cycle, with offsets 1, 1, 1 and loop sum $3 \neq 0$ — is precisely a configuration with **no** global section: locally decidable everywhere, globally frustrated. The triangle must be unfilled for this role: were a 2-cell attached along it, $d^1\rho$ would be nonzero (ρ evaluated on the cell's boundary sums to $3 \neq 0$), ρ would fail closedness, and the configuration would be excluded by Condition 1 rather than witnessing frustration. An unfilled triangle is automatically closed (no 2-cell for $d^1\rho$ to act on) but not exact, which is exactly the witness wanted.

The scoping already done by the arc yields the following status for candidate sources of GIP:

- **Energy conservation (BCB):** proven *not* to supply GIP. It gives closed, not exact — flatness, a strictly weaker condition. The Aharonov–Bohm situation is the standing illustration: a field closed on every accessible region, energy conserved, yet the potential not globally single-valued and the holonomy physically real.
- **Closure consistency at hubs:** proven not to supply GIP unaided. Locality does not force a global section; if it did, the T1 branch would be empty.
- **The data/dynamics cut:** no evident route. A committed datum can be a contextual comparison as readily as a global coordinate.
- **Internality / decomposition-independence (IA, PC):** orthogonal. They constrain how comparisons transform, not whether comparisons globalize.
- **Uniform Readout (UR):** the sole remaining candidate among the principles enumerated above. That UR is the *only* remaining candidate is conditional on this enumeration of VERSF principles being exhaustive — an enumeration-completeness claim inherited from the companion's scoping, not a proof. Whether UR forces a global section or only per-context sections is the logical-versus-strong contextuality distinction, and is the subject of this paper.

The burden is therefore located at a single axiom. Either GIP descends from UR — Branch A, $\kappa = 0$, sector empty by relabelling-triviality — or it does not, leaving Branch B live ($\kappa \neq 0$ possible, given availability).

---

4. The Equivalence Theorem

The first contribution is to make the UR $\Longrightarrow$ GIP question exact by characterizing UR-global — the property that UR forces a global section — through three independently meaningful conditions the arc had treated as separate vocabularies. The substantive content is the equivalence of those three characterizations (GIP, path-independent ledger, non-contextuality); the equivalence of UR-global with the first of them is essentially definitional, and the proof says so.

Theorem 4.1 (Uniform Readout entails vanishing charge iff the ledger is path-independent) (proven, given the dictionary of §1 and the cochain setting).

The following are equivalent:

1. **(UR-global)** Uniform Readout supplies, for every admissible offset ρ, a single global vertex labelling $a : V \to \mathbb{Z}_7$ such that readout at every site is consistent with a; equivalently UR forces a global section.
2. **(GIP)** $A \subseteq B^1$: every admissible ρ is exact, $\rho = d^0a$.
3. **(Path-independent ledger)** The committed readout record assigns to each site a value independent of the transport path along which it is read; equivalently, the ledger value transported around any closed loop returns to itself.
4. **(Non-contextuality)** The admissible comparison data admit a global section in the sheaf sense; the contextuality obstruction κ vanishes for every admissible configuration.

Consequently, UR $\Longrightarrow$ GIP holds iff readout consults a path-independent ledger iff the admissible data are non-contextual; and the negation of any one — UR forcing only per-context sections, $A \not\subseteq B^1$, a merely-comparable (path-dependent) ledger, or genuine contextuality — is equivalent to the negation of all, leaving Branch B live.

Proof.

(2 $\Longleftrightarrow$ 4). By the universal-coefficient theorem, $H^1(\Gamma\_vac; \mathbb{Z}_7) \cong \mathrm{Hom}(H_1(\Gamma\_vac; \mathbb{Z}), \mathbb{Z}_7) \oplus \mathrm{Ext}(H_0(\Gamma\_vac; \mathbb{Z}), \mathbb{Z}_7)$, and the Ext term vanishes because H_0 is free. So a class is determined by its pairing against *integral* H_1 — which captures both free generators and 7-torsion generators: a $\mathbb{Z}_7$-torsion class of integral H_1 pairs nontrivially through $\mathrm{Hom}(\mathbb{Z}_7, \mathbb{Z}_7) \cong \mathbb{Z}_7$. (This is the precise route, and it is why $\mathbb{Z}_7$ coefficients are the right detector for the §7 torsion; the naive "$\mathbb{F}_7$-linear duality on $H_1(\cdot; \mathbb{F}_7)$" would not state the integral-torsion content correctly.) Hence $\kappa = [\rho] = 0$ for all admissible ρ iff every admissible ρ pairs trivially with every integral homology class, iff every admissible ρ is a coboundary, iff $A \subseteq B^1$. Vanishing of the contextuality obstruction is by definition the existence of a global section. So (2) and (4) are the same statement in two vocabularies.

(1 $\Longleftrightarrow$ 2). A global vertex labelling $a : V \to \mathbb{Z}_7$ consistent with readout at every site is exactly a 0-cochain whose coboundary d^0a reproduces the offsets: consistency of readout at the two ends of

an edge u → v with a single global labelling means the recorded comparison ρ_{uv} equals $a_v - a_u$. Quantified over all admissible ρ, "UR supplies such an a for every admissible ρ" is precisely "every admissible ρ is in B^1," i.e. GIP. Conversely GIP supplies the labelling. So UR-global and GIP are equivalent — an equivalence that is essentially definitional, since UR-global is defined as "UR forces a global section" and GIP as "a global section exists for every admissible ρ." It carries exactly one substantive premise, made explicit here: that the recorded comparison ρ_{uv} *is* the difference $a_v - a_u$ of readouts in the global frame. If the offset were instead an independently committed quantity not defined as a difference of frame values, UR-frame could fix the site values a without thereby forcing $\rho = d^0a$, and the implication would fail. The equivalence holds precisely when offsets are readout differences — which the closure-connection definition of ρ as a pairwise comparison supplies.

(3 ⟺ 2). A ledger is path-independent iff transporting the recorded value around any closed loop returns it unchanged — iff the loop sum of the offsets vanishes on every cycle. Over $\mathbb{Z}_7$, "loop sum vanishes on every cycle" means $\langle \rho, z \rangle = 0$ for all $z \in Z_1$, which (since ρ is already closed by §2) means $[\rho]$ pairs trivially with all of H_1, hence $\kappa = 0$, hence $\rho \in B^1$. Quantified over admissible ρ this is $A \subseteq B^1$ = GIP. Conversely, exactness $\rho = d^0a$ makes the value a a global potential, manifestly path-independent. A merely-comparable ledger — one supporting pairwise comparisons with no global potential — is exactly a closed-but-not-exact ρ, path-dependent around an open cycle.

The four are therefore equivalent, and the equivalence of their negations follows immediately.

What the theorem delivers. It converts "does UR entail GIP" — a question about an axiom whose precise content must be read from the framework — into "does UR force a path-independent ledger," a question about the *physical nature of the committed record*. These are the same question, but the second is answerable by inspecting what commitment actually lays down, which §6 takes up. It also collapses three characterizations of UR-global — $A \subseteq B^1$, path-independent ledger, and non-contextuality — into one condition, with UR-global itself the object they characterize rather than a fourth vocabulary; and the negation cluster (path-dependent ledger / contextual / $A \not\subseteq B^1$ / per-context-only readout) is its single complement.

---

5. Decomposing Uniform Readout

Theorem 4.1 shows GIP is equivalent to UR carrying *global* force. But "uniform readout" is ambiguous between two strengths, and separating them shows exactly which component must do the work.

Uniform Readout, as scoped in the arc, requires that readout assign outcomes uniformly across admissible configurations. This decomposes into two independent claims:

- **(UR-comp) Comparison uniformity.** The readout of the *difference* between two sites is the same function of the offset wherever it is applied: the rule taking ρ_{uv} to a recorded comparison does not depend on which edge or context it acts in. This is a statement about the comparison component — about edges.

- **(UR-frame) Frame uniformity.** There is a single global frame relative to which every site's readout is referred: a global 0-cochain a such that each site's outcome is its value in that one frame. This is a statement about the frame component — about a global section.

Proposition 5.1 (Only the frame component can supply GIP) (proven). UR-comp does not entail GIP; UR-frame entails GIP; and UR-frame is strictly stronger than UR-comp.

Proof. Take the triangle as a *cochain witness* — three sites pairwise compared with offsets 1, 1, 1, the loop unfilled. Used here it need not be admissible; it serves only to separate UR-comp from UR-frame as logical conditions. It satisfies UR-comp: each edge applies the identical comparison rule, "the offset is 1," uniformly. Yet its loop sum is $3 \neq 0$ and it has no global section, so it is not exact — UR-comp holds while GIP fails. Hence UR-comp does not entail GIP. UR-frame, by definition, supplies the global labelling a; under the offset-as-difference premise of Theorem 4.1 (recorded offsets are differences of frame values) this gives $\rho = d^0a$, which is exactly GIP. And UR-frame $\Longrightarrow$ UR-comp (a single global frame induces a uniform comparison rule via differences), while the triangle shows the converse fails; so UR-frame is strictly stronger.

A separate remark on roles, to forestall a misreading against Condition 1. The triangle of this proof is a *logical* witness, used only to separate the two uniformities, and need not be admissible. The *occupancy* witness of §2 is a different object: to witness occupancy it must additionally be admissible — closed ($d^1\rho = 0$), and supported on a surviving cycle of a T1 complex (an unfilled free loop, or a degree-seven torsion loop per §7). Superficially similar configurations thus play two distinct roles; only the occupancy witness must satisfy admissibility, and an unfilled triangle qualifies as one (closed for lack of a filling 2-cell, non-exact for nonzero loop sum) while a filled triangle does not (it would violate closedness).

Consequence. The entire weight of Branch A rests on whether the framework's Uniform Readout axiom is UR-frame or merely UR-comp. If UR means only that the *comparison rule* is uniform (the same $\mathbb{Z}_7$ subtraction applied on every edge), it is satisfied by contextually frustrated data and does not force a global section — Branch B is live. If UR means that readout is referred to a *single global frame*, it is GIP outright — Branch A holds, and trivially so. The arc's "logical-versus-strong contextuality" fork is precisely the UR-comp / UR-frame fork: UR-comp permits logical/strong contextuality (per-context sections that need not glue); UR-frame forbids it (one global section).

The question "does UR entail GIP" is therefore not a deep derivation waiting to be done. It is the reading of a single axiom: which uniformity does Uniform Readout assert?

---

6. The Deciding Property and the Physical Fork

Proposition 5.1 reduces occupancy to one property of the readout axiom. Theorem 4.1 gives that property a physical face: does committing facts ever lay down a **global frame** (a committed 0-cochain $a : V \to \mathbb{Z}_7$, the path-independent ledger), or does it only ever commit **local comparisons** (committed 1-cochain offsets ρ on edges)? The two readings, stated without prejudice.

The Branch-A reading (single path-independent ledger; $\kappa = 0$). If the data side of the data/dynamics cut is a single committed ledger — one global record from which all readout is drawn — then that ledger *is* a global 0-cochain, readout is UR-frame, and GIP holds by Theorem 4.1. The sector is empty: every admissible offset is the difference of the ledger's values, a pure relabelling.

The honest weakness of this reading: it must establish that the committed data form a *single global frame*, which is exactly path-independence — assuming it begs the question. The data/dynamics cut, as scoped in the arc, separates committed data from reversible dynamics but does not by itself assert that the committed data are globally integrable; a committed datum can be a contextual comparison as readily as a global coordinate. So Branch A requires not merely that there is a ledger, but that the ledger is path-independent — and that is the very thing in dispute.

The Branch-B reading (only local committed comparisons; $\kappa \neq 0$ possible). If commitment only ever lays down *local* facts — and the measurement layer is explicit that a Bit is a discrete, local, irreversible threshold-crossing, a fact at a site or edge, not a global assignment — then the committed record is built entirely from local comparison offsets. Nothing in the act of committing a local fact establishes a global frame. The ledger is then exactly as integrable as the accumulated local facts force it to be — and on a T1 topology, the triangle countermodel shows local facts can be globally frustrated. Readout is UR-comp, not UR-frame; per-context sections need not glue; Branch B is live.

The honest weakness of this reading: "only local facts are ever committed" is a claim about the readout/commitment axiom that must be verified, not assumed. It is strongly suggested by the locality and irreversibility of the Bit, but the readout axiom could, in principle, include a global-frame-fixing component not visible in the Bit definition alone.

The two readings are not, however, symmetric in parsimony relative to the local generative chain, and §6.1 makes that asymmetry precise. It does not decide the fork — that is the test of §8 — but it identifies which branch carries the burden, and why.

6.1 Local Commitment Bias

The reduction of §6 leaves a single property of the readout and commitment layer: does commitment ever lay down a global frame, or only local comparisons? The reduction is exact. What remains is whether the existing ontology of the programme naturally favours one side of the fork. The answer is asymmetrical.

Throughout the VERSF construction the primitive objects are local. The generative chain is

Ticks → Fold → Bit → Fact → Record,

and each stage is defined through local distinction and local commitment. A Fold is a local closure event; a Bit is a local irreversible distinction; a Fact is a local commitment crossing a threshold; a Record is the accumulation of such local commitments. No primitive object in the construction is defined as a global assignment over the entire substrate. This observation does not prove Branch B, but it bears directly on the relative plausibility of the two branches.

Proposition 6.2 (Local Commitment Bias) (conditional on the local-ontology premise). Assume the primitive ontology of the programme consists exclusively of local commitment events. Then no primitive operation directly generates a global section $a : V \to \mathbb{Z}_7$. A path-independent ledger can therefore arise only as an *aggregate* — accumulated local offsets that happen to be jointly exact, $\rho = d^0 a$ for some a existing across the whole complex — and on a T1 topology such joint exactness is non-generic: generic closed local commitments have nonzero loop sum on the surviving cycles. So a path-independent ledger requires either an additional integrability mechanism relating independently committed regions, or a coincidence of joint exactness that the topology does not force.

This is not an independent theorem standing alongside the §3 closure-consistency negative; it is that negative restated at the level of ontological generation. The §3 result asks whether a *principle* forces GIP and finds that locality does not; Proposition 6.2 asks whether the commitment *ontology* generates a global frame and finds that it does not. The mathematical force is one fact — locality does not produce global integrability — viewed from two sides; the novelty here is only the burden-shift framing of Corollary 6.3, not fresh derivation. It is marked conditional, not "(proven)" as an independent result, for exactly this reason.

Argument. A global section is a 0-cochain assigning a value to every site simultaneously. A commitment event, by contrast, records only a local distinction: it contributes information to a bounded neighbourhood of the substrate but does not by itself specify values on distant sites. Accumulating local commitments may aggregate into a globally integrable structure, but only if the accumulated offsets happen to be jointly exact; on a T1 topology this is non-generic, since generic closed offsets have nonzero loop sum on the surviving cycles. Nothing in the primitive definitions of Tick, Fold, Bit, Fact, or Record supplies a mechanism forcing that joint exactness. Therefore the existence of a path-independent ledger is not a consequence of local commitment alone — it needs either an added integrability mechanism or a coincidence the topology does not enforce.

Corollary 6.3 (Burden Shift) (given Proposition 6.2). The two branches stand in an asymmetric relation to the local generative chain — but the asymmetry is one of parsimony, not of logical force or probability, and each branch carries its own premise. Branch A requires a *presence*: a global-frame-fixing mechanism, derived or posited, that converts locally committed comparisons into a single integrable frame. Branch B requires an *absence*: that no such mechanism exists — together with an available topology (T1). Branch B's premise of absence is precisely the local-ontology premise of Proposition 6.2, and is exactly the property the §8 test checks; so "Branch B requires nothing further" is true only once that premise is granted. The asymmetry is therefore this: given that every primitive in the chain is local, the absence is the *default* reading and the presence is the *addition* that must be exhibited. The burden of proof accordingly lies on any derivation of Branch A to identify the mechanism by which a global section is produced. It does not follow that Branch B is more likely — only that, relative to the established local primitives, it posits nothing beyond their absence plus T1, whereas Branch A posits a new global structure.

This burden cannot be discharged by relabelling. If Uniform Readout is itself posited at the axiom level as UR-frame, that posit *is* the additional global mechanism Corollary 6.3 demands — a global principle not supplied by the local generative chain, and an isolable assumption rather than a buried one. So the burden shift holds whether the integrability mechanism is derived or posited: either way Branch A must exhibit, and own, a global-frame-fixing principle that the local commitment ontology does not provide. The deciding question of §6 — does commitment lay down a global frame or only local comparisons? — is precisely whether that principle exists.

Remark 6.4 (the result must not be over-read). This does not establish Branch B. A globally integrable ledger could still emerge as a higher-level consequence of Uniform Readout or another admissibility principle; Proposition 6.2 establishes only that such a ledger is not supplied *directly* by the primitive commitment ontology. The local-ontology premise is reinforced by the ground state — the void has $H_1 = 0$ (no persistent distinctions, trivially a global section), and reality accumulates as fold networks depart from it by local commitment events — but, as Corollary 6.3 records, that premise is exactly what the §8 test checks, so the burden shift is parsimony relative to the existing primitives, not proof. The absence of a global-frame-fixing mechanism is not proof of Branch B; it identifies precisely what remains to be shown.

---

7. The Torsion Role of Seven

It is worth recording where the "7" genuinely enters, since the corrected object makes this precise and it is not where the bit-count draft placed it.

The seven appears in two related places, neither of them a phase period.

First, $\mathbb{Z}_7$ is the **coefficient group** of the offset cochains — the comparison alphabet, the $K = 7$ constraint count serving as the set of distinguishable pairwise offsets.

Second, and structurally independent of the first, the $K = 7$ count fixes an **attaching degree**. A closure constraint at a hub attaches a 2-cell to the complex; a *seven-fold* closure constraint attaches its 2-cell along a loop γ traversed seven times — a degree-seven attaching map, $\partial_2(\sigma) = 7\cdot\gamma$. In integral homology this makes $[\gamma]$ a **torsion class of order 7**: $7[\gamma] = 0$ while $[\gamma] \neq 0$, so $[\gamma]$ generates a $\mathbb{Z}_7$ summand of $H_1(\Gamma\_vac; \mathbb{Z})$. This is the concrete form of the torsion the survival criterion needs, and it gives $K = 7$ a second home — not only the coefficient alphabet but the attaching degree of the closure 2-cells.

Two consequences worth stating exactly.

(i) On such a torsion cell every $\mathbb{Z}_7$-cochain is automatically closed: $d^1\rho(\sigma) = \rho(\partial_2\sigma) = \rho(7\gamma) = 7\langle\rho, \gamma\rangle \equiv 0 \pmod 7$. So flatness imposes no constraint *there*. Flatness does its work on the ordinary degree-one cells — which fill loops and kill the corresponding free classes — while the degree-seven cells are exactly where surviving $\mathbb{Z}_7$ classes can live. This is the precise division of labour the corrected object requires: degree-one cells make flatness non-vacuous, degree-seven cells make the torsion possible.

(ii) A cochain with $\langle\rho, \gamma\rangle = 1$ on such a torsion loop is closed but not exact, hence represents $\kappa \neq 0$ — a clean occupancy witness supported on a torsion generator, complementary to the unfilled-free-loop witness of §3 and §5.

A point of precision the corrected object forces. The surviving class is a **torsion class in integral homology** $H_1(\Gamma\_vac; \mathbb{Z})$, of order 7 — *not* a "torsion class in $H_1(\Gamma\_vac; \mathbb{Z}_7)$," which would be a category error: with field coefficients $H_1(\cdot; \mathbb{Z}_7)$ is a $\mathbb{Z}_7$-vector space and has no torsion at all. The $\mathbb{Z}_7$ coefficients *detect* the integral 7-torsion — through the $\mathrm{Hom}(\mathbb{Z}_7, \mathbb{Z}_7) \cong \mathbb{Z}_7$ term of the universal-coefficient identification of §4 — but they are not where the torsion lives. The torsion is rationally invisible (a degree-seven loop bounds over $\mathbb{Q}$, since $7\gamma = \partial_2\sigma$ makes γ a boundary after tensoring with $\mathbb{Q}$), which is exactly why it is seen only with $\mathbb{Z}$ or $\mathbb{Z}_7$ coefficients and not in the free/rational homology.

So if Gate 3 is occupied, the charge's "sevenness" is the order of a torsion generator of the substrate's integral first homology — a structural property of how the closure constraints attach cells to the vacuum transport complex — not the period of a winding of the amplitude phase. The $K = 7$ count enters twice over: as the coefficient alphabet of the offsets, and as the attaching degree that creates the torsion. This is the honest home of the seven in the holonomy sector.

---

8. The Test: Well-Posed, Decidable Once UR's Quantifier Is Fixed

The reduction yields a single operational test, to be applied to the framework's Uniform Readout (and the associated commitment/ledger) axiom.

Test. Inspect the Uniform Readout axiom and determine which of the following it asserts:

(UR-frame) Readout refers every site's outcome to a single global frame — equivalently, commitment lays down (or readout presupposes) a global 0-cochain $a : V \to \mathbb{Z}_7$, a path-independent ledger.

(UR-comp) Readout applies a uniform comparison rule on each edge/context, but refers outcomes only to local per-context frames — commitment lays down only 1-cochain offsets on edges.

If UR-frame: GIP holds, $A \subseteq B^1$, $\kappa = 0$, **Gate 3 is unoccupied** (Branch A), and the closure sector is empty by relabelling-triviality.

If UR-comp: GIP fails, $A \not\subset B^1$ is permitted, and **Gate 3 is occupied** (Branch B) provided the topology supplies an open cycle (T1) — which the survival criterion then guarantees the witness occupies, since closedness forces any nonvanishing-holonomy offset onto an open cycle.

The test is well-posed because it asks only which uniformity the axiom asserts — a reading of the axiom's quantifier structure (does it quantify over a single global frame, or over per-context frames?), not a new derivation. It is *decidable* once the Uniform Readout axiom's quantifier is fixed: as §5–§6 note, the present statement of UR may not yet disambiguate UR-comp from UR-frame, and in that case the test returns not an answer but an instruction — which quantifier to sharpen. By Proposition 5.1 these are the only two options that bear on GIP, and by Theorem 4.1 they correspond exactly to the path-independent and merely-comparable ledger.

A sharper, equivalent form of the test, stated in terms of commitment: **Does any single commitment event ever lay down a global frame (a 0-cochain), or do all commitment events lay down only local comparisons (1-cochains)?** If the former is ever true and is forced by admissibility, Branch A; if commitment is exhaustively local, Branch B. This form is checkable directly against the Bit/commitment definition of the measurement and fact-production layers.

9. Status and Conclusion

Corrected. The closure charge is $\kappa \in H^1(\Gamma\_vac; \mathbb{Z}_7)$, a contextuality obstruction with $\mathbb{Z}_7$ as cochain coefficient group, not a winding of the amplitude phase. The bit-count framing and its connectedness no-go, though internally correct, concern an object the framework does not use, and are superseded by the present treatment.

Proven (this paper).

- Theorem 4.1: UR entails $\kappa = 0$ for all admissible configurations iff readout consults a path-independent ledger iff the admissible data admit a global section (non-contextuality). The arc's vocabularies — UR-global, $A \subseteq B^1$, path-independent ledger, non-contextuality — are one condition.
- Proposition 5.1: Uniform Readout splits into a comparison component (UR-comp) and a frame component (UR-frame); only UR-frame entails GIP, and UR-frame is strictly stronger; the logical-versus-strong contextuality fork is exactly this split.

Reduced to one well-posed input. Occupancy now rests entirely on which uniformity the Uniform Readout axiom asserts — equivalently, on whether commitment ever lays down a global frame or only local comparisons. The test of §8 is the operational form: well-posed now, decidable once UR's quantifier is fixed.

Open (the single remaining input). Whether the framework's Uniform Readout / commitment axiom is UR-frame (→ Branch A, $\kappa = 0$, sector empty) or UR-comp (→ Branch B live, $\kappa \neq 0$ possible on a T1 topology). The Local Commitment Bias (Proposition 6.2) and the Burden Shift (Corollary 6.3) record a parsimony asymmetry relative to the local generative chain: Branch A requires the *presence* of a global-frame-fixing mechanism the commitment chain does not supply, while Branch B requires its *absence* together with an available topology (T1). Both carry a premise — Branch B's being the local-ontology premise the §8 test checks — so the asymmetry is one of parsimony relative to the established primitives, not of logical force or probability. It locates the burden on Branch A and is not to be read as support for Branch B.

Not claimed. This paper does not prove Gate 3 is occupied, does not prove it is empty, and does not treat the weakening of Branch A as support for Branch B. It proves that the entire occupancy question has converged onto one well-posed property of one axiom — decidable once UR's quantifier is fixed — supplies the equivalence that gives that property a physical face (path-independence of the ledger), and hands forward the test.

The shape of the terminus is now clean. Three of the arc's conditions are discharged: conservation gives flatness, the topology paper gives availability, the survival criterion fuses them. What remains is a single reading of a single principle. The closure charge survives if and only if reading out the substrate is consulting a record that depends on the path — a ledger of local comparisons that no global frame reconciles — and it vanishes if and only if readout consults one path-independent global frame. Whether the substrate commits a global frame or only local comparisons is the whole of Gate 3, and it is now a checkable property of the readout axiom rather than an open structural mystery.

The closure charge is a failure of local comparisons to glue into a global frame — a contextuality obstruction valued in $\mathbb{Z}_7$. It survives exactly when reading the substrate means following a path-dependent ledger of local facts, and vanishes exactly when one global ledger reconciles them. The arc has reduced to one question: does committing a fact lay down a global frame, or only a local comparison?