The Continuum Lift of the Gate-3 Closure Sector

From κ ∈ $H^1(\Gamma\_vac; \mathbb{Z}_7)$ to Closure-Memory Curvature Ω

Keith Taylor — VERSF Theoretical Physics Programme

General Reader Summary

The VERSF programme has, in two separate places, found what looks like *leftover* structure — information that survives after the substrate has finished settling locally. One place is the Gate-3 work, which describes this leftover as a discrete, sevenfold ($\mathbb{Z}_7$) "residue" living on the vacuum closure complex. The other is the Substrate Response Principle paper, which describes how committed information is read back from the substrate, and finds that the readout can carry an intrinsic "twist" (a curvature). The two descriptions look different: one is discrete and global, the other continuous and local.

This paper asks whether they are secretly the same thing. It does **not** claim they are. What it does is pin down the exact condition under which they *could* be identified — and shows that condition is sharp and non-generic: the continuous twist would have to be *quantized* into exactly seven values when carried around the closed loops of the vacuum complex, rather than ranging freely as a continuous twist normally would. The simplest way that can happen is if the twist vanishes locally but leaves a sevenfold remainder on the loops; but it can also happen, as in ordinary gauge physics, with a genuinely non-zero twist whose total around each loop still lands on one of seven values. Either way, the decisive thing is that the twist be *quantized to seven*, not that it vanish. If the substrate's readout has that quantized form, the two leftovers are one object in two notations. If the twist instead ranges continuously, they are genuinely different, and the programme has two separate open problems rather than one — a continuous twist cannot be squeezed into seven values without extra information the twist itself does not carry, which is the one negative result the paper proves outright.

The paper's contribution is to convert a vague "are these the same?" into a single yes-or-no condition that can be checked, and to determine how much of the answer the existing corpus already fixes (some of it) versus leaves open (the decisive part). It is a problem-sharpening paper, not a resolution.

Epistemic markers: ***(established)*** *for results inherited from prior VERSF papers;* ***(constraint)*** *for conditions the corpus forces on any candidate lift;* ***(conjecture)*** *for proposed identifications not derived here;* ***(open)*** *for what remains undecided.*

Central Finding

The most consequential result of this paper is not the lift itself but the discovery of what governs it. The entire strength of the proposed identification between the reverse-map curvature Ω and the Gate-3 closure class κ turns on a single, finite, mechanically computable invariant of the vacuum complex:

$\mathrm{Tor}(H_1(\Gamma\_vac; \mathbb{Z}))$ — the torsion subgroup of the first integral homology.

The two possible outcomes are not equally interesting, and the difference between them is the difference between the paper having found real structure and having found a coincidence of notation:

- **If the $\mathbb{Z}_7$ is only a coefficient choice** — i.e. $H_1(\Gamma\_vac; \mathbb{Z})$ is torsion-free and $\mathbb{Z}_7$ enters merely as the chosen coefficient group in $H^1(\Gamma\_vac; \mathbb{Z}_7)$ — then the sevenfoldness is not in the space. One could equally write $H^1(\Gamma\_vac; \mathbb{Z}_5)$ or $H^1(\Gamma\_vac; \mathbb{Z}_{11})$ on the same complex; the seven is a label on the reverse-map side, contributing nothing from the topology. The connection gains nothing from the Gate-3 sector, and the identification is **weak** — two structures sharing a coefficient convention.
- **If $\mathrm{Tor}(H_1(\Gamma\_vac; \mathbb{Z})) = \mathbb{Z}_7$** — genuine sevenfold torsion in the integral homology — then the seven is *in the space*, surviving any choice of coefficients. The connection inherits it through the torsion cycles (the mechanism of §6.2, made concrete in the worked example of §4.4), the two sevenfoldnesses have a common topological cause, and the identification is **genuine**. This is the branch in which the paper has found real structure rather than a notational coincidence — the branch in which it becomes interesting.

A precise caveat keeps this honest. The torsion computation decides whether the sevenfoldness is *genuine or notational* — and that is the more important of the paper's two open facts, because it is mechanically checkable now and it determines whether the identification is worth pursuing at all. But the favourable (torsion) outcome **earns** the paper's ambition; it does not by itself **cash** it. Even with $\mathrm{Tor}(H_1) = \mathbb{Z}_7$, completing the lift still requires the continuum limit to preserve that torsion in the actual reverse-map holonomy (§6.3, the second open fact). So: the coefficient outcome would collapse the paper's ambition; the torsion outcome establishes that the structure is real and the lift worth pursuing; and the continuum limit is what would finish it. The single most valuable next action in the entire programme-thread this paper belongs to is therefore to compute $\mathrm{Tor}(H_1(\Gamma\_vac; \mathbb{Z}))$ directly — a finite calculation on a complex already in hand.

Table of Contents

---

1. Introduction

Two apparently separate open questions have emerged within the programme, and both concern structure that survives closure.

The Gate-3 line asks whether the vacuum closure complex possesses a non-trivial residual closure class

$\kappa \in H^1(\Gamma\_vac; \mathbb{Z}_7)$.

The reverse-map line, developed in the Substrate Response Principle paper, asks whether reconstruction of committed information is governed by a flat or a curved connection — equivalently, whether the closure-memory curvature

$\Omega = d\Gamma$

vanishes.

These have been treated independently. Yet both describe residual structure surviving closure: both concern information remaining after local completion, and both appear as obstructions to global triviality. This raises the possibility that they are not independent questions but two notations for one substrate degree of freedom.

The purpose of this paper is to determine whether the two sectors admit a common interpretation — and, crucially, to do so without assuming the answer. We state the identification as a hypothesis with two horns, derive the exact condition that would decide between them, and establish how much of that condition the corpus already settles. We do **not** claim to complete the identification. We claim only to make it a sharp, checkable question rather than a suggestive analogy.

It is worth naming the paper's deepest vulnerability at the outset rather than letting the reader find it. The reverse-map connection Γ was introduced, in the Substrate Response Principle paper, *by analogy* to the gauge connection A_μ. A skeptic may therefore object that Ω is only analogically a curvature, that κ is a genuine cohomological object while Ω is a heuristic one, and that any "identification" of the two is analogy stacked on analogy rather than a relation between comparably real structures. We take this objection seriously and address it directly in §6.2: the defense is that, *if* both Γ and the Gate-3 sector live over the same K = 7 vacuum complex, their relationship is structural — a shared substrate — rather than analogical, and the sevenfoldness of each has a common cause. That defense is conditional on a homological fact about Γ_{vac} that we do not assume and flag explicitly (§6.2). The honest position is therefore: the objection is real, the defense is structural rather than analogical, and the defense's premise is a named open computation — not a hope.

---

2. The Two Objects

2.1 The Gate-3 Closure Class (established, Gate-3 line)

The Gate-3 programme defines the residual closure sector by the cohomology class

$\kappa \in H^1(\Gamma_{vac}; \mathbb{Z}_7)$.

Operationally, κ measures the failure of closure transport around admissible loops to reduce globally to the trivial class. Its essential features:

- κ is **discrete** — valued in $\mathbb{Z}_7$, the integers mod 7.
- κ is **topological** — a cohomology class, invariant under local deformation.
- κ is **global** — it survives only if non-trivial first cohomology of Γ_{vac} remains; it is not a local field but a global invariant.

The Gate-3 line attributes the $\mathbb{Z}_7$ valuation to the K = 7 simplicial structure of the vacuum complex. Whether this means $\mathbb{Z}_7$ is intrinsic to the homology of Γ_{vac} (as torsion) or only the chosen coefficient group is a distinction that turns out to be decisive for the identification; we defer it to §6.2, where it becomes load-bearing.

2.2 Closure-Memory Curvature (established, SRP §9.1)

The reverse-map programme introduces a connection Γ governing reconstruction transport, with curvature

$\Omega = d\Gamma$ (abelian), $\Omega = d\Gamma + \Gamma\wedge\Gamma$ (non-abelian),

and holonomy around a loop ℓ

$H(\ell) = \oint_\ell \Gamma$ (abelian), $H(\ell) = \mathcal{P}\exp(\oint_\ell \Gamma)$ (non-abelian).

In the readout fork of the SRP paper, Branch A corresponds to $H(\ell)$ trivial on all contractible loops (Γ flat), Branch B to non-trivial holonomy for some loop (Γ curved). Its essential features:

- Ω is **continuous** — a curvature 2-form.
- Ω is **local** — it assigns curvature point-by-point across the substrate.
- Ω is **Lie-algebra-valued** — its holonomy lives in a continuous group, via the gauge-connection analogy from which Γ is built.

2.3 Why they appear to belong to different categories

Laid side by side, the two objects differ in every classifying attribute:

Attribute	κ (Gate-3)	Ω (reverse-map)
Valuation	discrete, $\mathbb{Z}_7$	continuous, Lie-algebra
Locality	global invariant	local field
Type	cohomology class	curvature 2-form
Origin	K = 7 complex topology	reconstruction connection

On their face these are objects of different mathematical kinds, and one cannot simply *equate* them. The question of this paper is whether the appearance of difference is deceptive — whether κ is what Ω's holonomy *becomes* after projection into the discrete K = 7 setting, or whether the two are genuinely independent.

---

3. The Continuum-Lift Hypothesis — and Its Two Horns

We state the proposed identification as a hypothesis, explicitly **not** as a result, and we attach to it the alternative it must be tested against.

Continuum-Lift Hypothesis (conjecture)

The Gate-3 closure class κ is the topological residue generated by closure-memory curvature Ω — that is, κ is what the continuous holonomy $H(\ell)$ reduces to after projection into the K = 7 closure algebra:

$\Omega \longrightarrow H(\ell) \longrightarrow \kappa$.

Under this reading Ω is the local closure-memory curvature, H(ℓ) its integrated holonomy around admissible loops, and κ the surviving topological residue after admissibility projection — and the Gate-3 and reverse-map sectors are one substrate channel in two notations.

The Distinctness Alternative (the other horn)

Ω and κ are genuinely independent structures. The continuous curvature of the reconstruction connection and the discrete $\mathbb{Z}_7$ residue of the vacuum complex are different substrate degrees of freedom that happen to share the informal description "structure surviving closure." Under this reading the SRP paper's "closure-memory" language carries no commitment to the Gate-3 sector, and the two open problems remain two.

These horns are mutually exclusive, and a single sharp condition decides between them, derived next. We do not presuppose which holds; we derive the selecting condition and ask what the corpus says about it.

---

4. What a Lift Would Require — The Reduction Map

4.1 The valuation problem

The central obstruction to any lift is the valuation mismatch of §2.3. H(ℓ) is valued in a continuous group; κ is valued in $\mathbb{Z}_7$. A map

$H(\ell) \longrightarrow \kappa \in \mathbb{Z}_7$

cannot be an identity — it must be a **projection** that collapses the continuous holonomy onto a finite group. Such a projection exists only if the continuous holonomy, evaluated on the admissible loops of Γ_vac, already takes values in (or surjects onto) a $\mathbb{Z}_7$ subgroup of the holonomy group. A continuous holonomy with generic Lie-algebra value does **not** reduce to $\mathbb{Z}_7$: its image is a continuum, not seven points. So the lift is not generic. It requires the holonomy to be quantised, on the relevant loops, into exactly the sevenfold set.

4.2 The flat-mod-$\mathbb{Z}_7$ condition (constraint)

There is a precise way for a continuous connection to have $\mathbb{Z}_7$-valued holonomy: it must be **flat modulo** $\mathbb{Z}_7$. That is,

- the curvature vanishes in the continuous sense, $\Omega = 0$ on contractible regions (so there is no continuous holonomy generated by local curvature), **and**
- the connection carries a non-trivial *residual* holonomy on the non-contractible loops of Γ_vac, valued in the $\mathbb{Z}_7$ subgroup determined by the K = 7 structure.

Equivalently, Γ is a flat connection whose holonomy representation $\pi_1(\Gamma\_vac) \to G$ factors through $\mathbb{Z}_7 \subset G$. The holonomy is then a homomorphism from loops to $\mathbb{Z}_7$ — exactly the data of a class in $H^1(\Gamma\_vac; \mathbb{Z}_7)$. Under this and only this condition does $H(\ell)$ deliver a well-defined element of $\mathbb{Z}_7$ for each loop class, and hence a candidate κ.

We state this as the **flat-mod-$\mathbb{Z}_7$ condition**:

$\Omega = 0$ (continuous curvature) and $\mathrm{hol}(\Gamma) : \pi_1(\Gamma\_vac) \to \mathbb{Z}_7$ non-trivial.

This is the cleanest and most direct route to the lift: it is *sufficient* for a canonical identification, as §4.3 shows. We are careful **not** to claim it is necessary. A connection can, in general, carry non-zero local curvature and still produce quantized global holonomy — in ordinary gauge theory this is routine, where a curved connection yields discrete global charges because the curvature integrates to a quantized flux over the relevant cycles. So there may exist curved Ω whose holonomy is nonetheless $\mathbb{Z}_7$-valued on the admissible loops, supplying a κ without Ω vanishing. We have not excluded that route and do not claim flat-mod-$\mathbb{Z}_7$ is the only one. What we claim is sharper and weaker: flat-mod-$\mathbb{Z}_7$ is *one* condition under which the lift is canonical and forced, and it is the simplest such condition. The genuinely necessary condition is weaker still, and is isolated in the Proposition of §4.3.

4.3 The reduction map, stated conditionally (conditional)

If Γ satisfies the flat-mod-$\mathbb{Z}_7$ condition, the reduction map is forced and explicit:

$\kappa([\ell]) := \mathrm{hol}(\Gamma, \ell) \in \mathbb{Z}_7$, for $[\ell] \in \pi_1(\Gamma\_vac)$,

which is well-defined on homotopy classes precisely because $\Omega = 0$ makes the holonomy deformation-invariant, and $\mathbb{Z}_7$-valued because the holonomy representation factors through $\mathbb{Z}_7$. This assignment is a homomorphism, hence an element of $\mathrm{Hom}(\pi_1(\Gamma\_vac), \mathbb{Z}_7) \cong H^1(\Gamma\_vac; \mathbb{Z}_7)$, which is κ.

We emphasise what this does and does not establish. It establishes that *under the flat-mod-$\mathbb{Z}_7$ condition*, Ω's data reduces to exactly κ — the lift is then not merely possible but canonical. It does **not** establish that Γ satisfies the condition, nor that the condition is necessary. A flat Ω with trivial residual holonomy yields the trivial class ($\kappa = 0$); a curved Ω may or may not lift, depending on whether its holonomy is quantized. The reduction map is conditional, and the condition is sufficient, not necessary.

Proposition 4.1 — No lift from generically curved Ω (proven)

If Ω is continuously curved with generic holonomy — that is, if the holonomy map hol(Γ) : π_1(Γ_vac) → G has image a continuum rather than a finite subgroup — then no canonical class κ ∈ H^1(Γ_vac; $\mathbb{Z}_7$) can be defined from Ω alone.

Proof. A class in H^1(Γ_vac; $\mathbb{Z}_7$) assigns to each loop one of seven values, and is a homomorphism π_1(Γ_vac) → $\mathbb{Z}_7$. If the image of hol(Γ) is a continuum, then hol(Γ) takes more than seven values on loops, and no homomorphism to $\mathbb{Z}_7$ can agree with it: any map collapsing a continuous image onto seven points must identify holonomies that hol(Γ) distinguishes, and the choice of which to identify is not supplied by Ω. Defining κ would therefore require additional quantization data — a specified projection G → $\mathbb{Z}_7$ — not contained in the connection. Hence no κ is *canonically* determined by a generically-curved Ω. (The deeper reason is structural: a connected holonomy group admits no non-trivial continuous homomorphism to a discrete group, so any $\mathbb{Z}_7$-valued reduction of a continuous-image holonomy is necessarily discontinuous — i.e. requires a non-canonical choice of extra data — rather than being induced by the connection.)

This is the honest content of the negative side. Note precisely what it excludes and what it leaves open:

- It **excludes** identifying κ with a generically-curved Ω — one whose holonomy genuinely ranges over a continuum. This is a real result: the casual assumption that any continuous curvature projects onto a discrete class is false, because the projection requires quantization data the curvature does not carry.
- It does **not** exclude a connection that is flat on the accessible region but carries quantized flux confined to inaccessible cycles (holes / H_1 torsion) — the Aharonov–Bohm/Dirac pattern of Position 3 (§7). There the holonomy *is* deformation-invariant on the accessible loops and lands in $\mathbb{Z}_7$, so a κ exists. The excluded case is specifically free curvature on the region the loops sweep; confined flux off that region is not excluded. Flat-everywhere (Position 2) and confined-flux (Position 3) are the two admissible lifting cases, differing only in where the sevenfold structure sits.

So the genuinely necessary condition for the lift is **$\mathbb{Z}_7$-quantized, deformation-invariant holonomy on the admissible loops** — which requires flatness on the accessible region (so that holonomy is a function of homotopy class at all) together with quantization into $\mathbb{Z}_7$. Vanishing curvature *everywhere* is sufficient but not necessary: the sevenfold structure may instead be carried by quantized flux confined to inaccessible cycles, leaving the accessible region flat (Position 3). What is necessary is flatness where the loops sweep, not flatness everywhere. The lift exists iff the holonomy is $\mathbb{Z}_7$-quantized and deformation-invariant on the accessible loops; flat-everywhere is the simplest such case, not the only one.

4.4 The lifting class is non-empty — a worked example

Proposition 4.1 is a no-go: it excludes the generic case. On its own it leaves a hostile question open — *has the favourable case been shown realisable at all, or only excluded its complement?* We answer that here by exhibiting, concretely, a complex and a connection that lift to a non-trivial κ. This establishes that the lifting class (Positions 2–3) is non-empty, which a no-go alone cannot do.

The complex. Take the mod-7 Moore space $M(\mathbb{Z}_7, 1)$: start with a circle S^1 (one 0-cell, one 1-cell e generating $\pi_1 = \mathbb{Z}$), and attach a single 2-cell along the degree-7 map $e^7 : \partial D^2 \to S^1$. Call the result X. Its cellular chain complex is

$\mathbb{Z} \to(\times 7)\to \mathbb{Z} \to \mathbb{Z}$, $(\partial_2 = \times 7, \partial_1 = 0)$,

giving $H_1(X; \mathbb{Z}) = \mathbb{Z} / 7\mathbb{Z} = \mathbb{Z}_7$. So X has exactly the homological $\mathbb{Z}_7$ torsion the favourable reading of §6.2 requires — here by construction, exhibiting the structure rather than assuming it of Γ_vac.

The connection. Put the flat connection on X whose holonomy representation sends the generator [e] to a generator of $\mathbb{Z}_7 \subset U(1)$ — explicitly, $\mathrm{hol}([e]) = \zeta := \exp(2\pi i/7)$. This is well-defined and flat: the 2-cell is attached along e^7, and consistency requires the holonomy around its boundary to be trivial, $\mathrm{hol}(e^7) = \mathrm{hol}([e])^7 = \zeta^7 = 1$, which holds. The attaching condition is exactly what *forces* the holonomy into the seventh roots of unity — the $\mathbb{Z}_7$ is not imposed by hand but required by the degree-7 attachment, which is the toy analogue of the "7·[loop] bounds ⟹ hol^7 trivial" mechanism of §6.2.

Which group the mechanism lives in. A point of precision, because it is exactly what distinguishes the toy from the general case. The holonomy here is **abelian** (valued in $\mathbb{Z}_7 \subset U(1)$), and an abelian-valued holonomy factors through the abelianisation: $\mathrm{Hom}(\pi_1, U(1)) = \mathrm{Hom}(H_1, U(1))$, since U(1) sees only π_1^ab = H_1. So although we wrote the generator as a loop, the connection genuinely depends only on its class in H_1, and the operative group is H_1, not π_1. For the Moore space this is harmless because $\pi_1(X) = \mathbb{Z}_7 = H_1(X)$ coincide ($\mathbb{Z}_7$ is abelian, so equals its own abelianisation). But the coincidence is not generic, and naming it matters: §6.2's mechanism is stated in H_1 ("H_1(Γ_vac; $\mathbb{Z}$) carries $\mathbb{Z}_7$ torsion"), the "7·[loop] bounds" condition is an H_1 statement (bounding = trivial in homology), and the example realises *that* H_1-torsion mechanism specifically. In the real Γ_vac, where π_1 may be non-abelian and $\pi_1 \neq H_1$, the abelian-holonomy reduction to H_1 is what keeps the argument well-posed — the lift's $\mathbb{Z}_7$ lives in $H_1 = \pi_1$^ab regardless, because the holonomy is abelian. The toy is honest about this: it works in H_1, which is where §6.2 claims the mechanism lives.

The lift. The holonomy is deformation-invariant (the connection is flat, $\Omega = 0$) and $\mathbb{Z}_7$-valued (it lands in $\langle\zeta\rangle \cong \mathbb{Z}_7$). Factoring through H_1 as just noted, it defines via the reduction map of §4.3

κ = (the class sending the H_1-generator to $\zeta \mapsto 1$) $\in \mathrm{Hom}(H_1(X), \mathbb{Z}_7) \cong H^1(X; \mathbb{Z}_7)$,

a non-trivial element. So $\kappa \neq 0$: the connection lifts to a non-trivial closure class.

What this establishes. Positions 2–3 are not vacuous. There exists a complex with $\mathbb{Z}_7$ homological torsion carrying a flat connection whose holonomy is exactly a non-trivial $\mathbb{Z}_7$ class — the favourable case is realisable, not merely un-excluded. Together with Proposition 4.1 (the generic case fails) this gives both sides of the dichotomy: the lift fails for generic curvature and succeeds for at least one concrete flat-mod-$\mathbb{Z}_7$ connection. What the example does **not** do is show Γ_vac *is* such a complex — that is the deferred homological fact of §6.2, and the example is deliberately a generic toy (the Moore space), not a claim about the vacuum complex. Its role is existence: to convert "the lifting condition might be satisfiable" into "here is something that satisfies it," and to give a template for the positive case — checking whether $H_1(\Gamma\_vac; \mathbb{Z})$ carries the $\mathbb{Z}_7$ factor the Moore space has by construction.

---

5. What the Corpus Fixes, Forces, and Leaves Open

We separate what the existing programme already determines about the flat-mod-$\mathbb{Z}_7$ condition from what it leaves open.

Fixed by the corpus. The $\mathbb{Z}_7$ target is not a free choice: it is forced by the K = 7 simplicial structure of Γ_vac, established in the Gate-3 line. If a lift exists at all, its discrete group is $\mathbb{Z}_7$ and nothing else — the sevenfoldness is inherited, not fitted. Likewise, the *form* of the reduction map (holonomy of a flat connection as a cohomology class) is standard once flatness-mod-$\mathbb{Z}_7$ is granted; there is no freedom in how the lift would work, only in whether its precondition holds.

Forced negatively. The corpus already excludes the *generic* case. Because κ is $\mathbb{Z}_7$-valued, a generically-curved Ω — one with continuous-valued holonomy — is incompatible with identification (Proposition 4.1): its holonomy cannot equal a $\mathbb{Z}_7$ class without extra quantization data Ω does not supply. So the corpus forces a dichotomy at the level of holonomy: either the holonomy is $\mathbb{Z}_7$-quantized and deformation-invariant on the admissible loops (lift exists — by a globally flat connection, or by flux confined off the accessible region) or curvature is free on the accessible region and the holonomy continuous-valued (lift fails, sectors distinct). There is no identification of a *generically* curved Ω with κ. The lifting condition is quantization of the holonomy, not flatness specifically; flatness on the accessible region is required (§7), but the sevenfold structure may sit either in a flat connection's topology or in confined flux, so the flat-everywhere case is the simplest quantized case, not the only admissible one.

Left open. Whether Γ actually satisfies the flat-mod-$\mathbb{Z}_7$ condition is **not** fixed by the corpus. This is exactly the SRP paper's open reverse-map question (its Q4, the value of Ω) refracted through the discrete structure: deciding the lift requires deciding whether Ω = 0 with residual $\mathbb{Z}_7$ holonomy, which is a substrate-level fact about the reconstruction connection that neither paper computes. The continuum-limit construction that would relate the smooth connection Γ to the simplicial complex Γ_vac is itself unbuilt (the recurring continuum-limit gap of the programme), and the lift question is downstream of it: one cannot rigorously project a continuous holonomy onto Γ_vac's loops without that limit.

The honest status is therefore: the lift's *machinery* is fixed ($\mathbb{Z}_7$ target, holonomy-as-cohomology reduction), the *generic case* is excluded, and the *decisive precondition* (flat-mod-$\mathbb{Z}_7$) is open and entangled with the continuum-limit problem.

---

6. Why Seven? — The Source of the $\mathbb{Z}_7$ on the Connection Side

There is a circularity in the argument so far that must be confronted rather than left implicit. The paper has repeatedly invoked "$\mathbb{Z}_7$-quantized holonomy" as the lifting condition — but the *reason* given for the sevenfoldness has been: because the target κ is $\mathbb{Z}_7$-valued, so if a lift exists its holonomy must land in $\mathbb{Z}_7$. That is true but circular: it derives the connection's sevenfoldness *from the assumption that the lift exists*, rather than from anything intrinsic to the reverse-map connection. A skeptical reader is entitled to ask the question this section addresses directly:

Why should the reverse-map connection know anything about sevenfoldness at all?

If the only answer is "because the target is $\mathbb{Z}_7$," then the $\mathbb{Z}_7$ on the connection side is *donated* by the Gate-3 side, and the identification is correspondingly weaker — less an identification of two independently sevenfold structures than a relabelling that imports κ's $\mathbb{Z}_7$ onto Ω by assumption. To be a genuine identification, the sevenfoldness must arise on the connection side from something that does **not** presuppose the lift.

6.1 The asymmetry to be repaired

The Gate-3 side is the one with a candidate account of its sevenfoldness, and §6.2 will make precise what that account must establish. For now: κ takes values in $\mathbb{Z}_7$, and the Gate-3 line attributes this to the K = 7 simplicial structure of the vacuum complex. We note immediately that "$\kappa \in H^1(\Gamma\_vac; \mathbb{Z}_7)$ is $\mathbb{Z}_7$-valued" is, by itself, only the statement that $\mathbb{Z}_7$ was chosen as the coefficient group — true of any space and not yet a fingerprint of K = 7. The substantive claim, which §6.2 isolates, is the stronger one that K = 7 puts $\mathbb{Z}_7$ into the *homology* of Γ_vac (as torsion), not merely into the coefficients. The reverse-map side has no independent account at all: the connection Γ was introduced (in the SRP paper) as a transport rule for reconstruction, with no sevenfoldness built in, and its holonomy group is a continuous group with no a priori reason to break to $\mathbb{Z}_7$. The asymmetry is therefore: κ's $\mathbb{Z}_7$ has a candidate explanation (pending the homology-vs-coefficient question), Ω's putative $\mathbb{Z}_7$ has none. Repairing it requires a reason, internal to the connection, for $\mathbb{Z}_7$ to appear.

6.2 The one non-circular route — shared K = 7 substrate

There is exactly one route by which the connection could acquire $\mathbb{Z}_7$ holonomy without presupposing the lift, and it turns on a fact the two objects share: **both live over the same K = 7 structure.** Γ_vac *is* a K = 7 simplicial complex. If the reverse-map connection Γ is genuinely defined on that same complex — transport along the edges and loops of the K = 7 vacuum complex, rather than along a smooth manifold that merely approximates it — then its holonomy is holonomy around the loops of a sevenfold complex, and the available values are constrained by the complex's combinatorial structure.

But whether this *forces* connection-side $\mathbb{Z}_7$ depends on a distinction that must be stated explicitly, because the route stands or falls on it — and the paper cannot adjudicate it from outside the Gate-3 line. The question is **what kind of object the Gate-3 $\mathbb{Z}_7$ is**:

- **If $\mathbb{Z}_7$ is merely a coefficient choice** — that is, $\kappa \in H^1(\Gamma\_vac; \mathbb{Z}_7) = \mathrm{Hom}(H_1(\Gamma\_vac), \mathbb{Z}_7)$ with $\mathbb{Z}_7$ selected as the coefficient group — then "shared K = 7 complex" does **not** force connection-side $\mathbb{Z}_7$. A G-connection on Γ_vac has G-valued holonomy regardless of the complex's combinatorics; its factoring through $\mathbb{Z}_7$ is a separate fact needing separate argument. Note $H^1(S^1; \mathbb{Z}_7) = \mathbb{Z}_7$ as well — the $\mathbb{Z}_7$ in the coefficients is generic to the coefficient choice, not a fingerprint of K = 7. Under this reading §6.2 does **not** escape the circularity; it collapses into the weaker parsimony argument of §6.4.
- **If $\mathbb{Z}_7$ is intrinsic to the topology** — that is, $H_1(\Gamma\_vac; \mathbb{Z})$ carries genuine $\mathbb{Z}_7$ *torsion* (equivalently $\pi_1(\Gamma\_vac)^{ab}$ has a $\mathbb{Z}_7$ factor), traceable to the K = 7 structure — then §6.2 works as a genuine result. A connection on Γ_vac inherits $\mathbb{Z}_7$ holonomy on the torsion cycles: since 7·[loop] bounds, $\mathrm{hol(loop)}^7$ encloses trivial flux, forcing hol(loop) into the seventh roots — sevenfold holonomy *from the connection side*, with the sevenfoldness genuinely shared rather than donated. This reading also rehabilitates Position 3 cleanly, because the torsion cycles are exactly the "holes" through which confined flux threads (§7).

The paper does not assume which reading holds — the distinction is a yes/no fact about Γ_vac that resolves by direct computation, and we state both branches so the result follows the moment the computation is in hand:

- **Computation: $H_1(\Gamma\_vac; \mathbb{Z})$ carries $\mathbb{Z}_7$ torsion.** Then §6.2 is a genuine escape from the circularity — connection-side sevenfoldness is forced by the torsion cycles, shared with κ rather than donated — and the worked example of §4.4 shows this branch is not vacuous (it exhibits exactly such a complex and connection). The non-circularity claim stands; cite the Gate-3 homology computation.
- **Computation: $H_1(\Gamma\_vac; \mathbb{Z})$ is torsion-free and $\mathbb{Z}_7$ is only the coefficient group.** Then §6.2's non-circularity claim must be withdrawn, and the case for connection-side sevenfoldness rests on the parsimony argument of §6.4 alone.

So the strength of §6 turns on a single, named, checkable quantity: the torsion subgroup of $H_1(\Gamma\_vac; \mathbb{Z})$. This is itself precisely what Q7 demanded — the two notations cannot be reconciled without knowing whether the Gate-3 $\mathbb{Z}_7$ lives in the homology or only in the coefficients — and it reduces the entire non-circularity question to one homology computation on the vacuum complex. We flag it as the priority cross-corpus calculation: it is the largest single increase in what this paper establishes, and it is internal to the Gate-3 construction rather than dependent on the continuum limit.

6.3 Why even the torsion reading is not yet established

Suppose the favourable (torsion) reading holds. The route is then *available* but still not *established*, for the recurring reason. It requires that Γ genuinely descend to (or be defined on) the K = 7 simplicial complex Γ_vac. But the SRP connection Γ is a *continuous* object on a smooth setting; whether it descends to the discrete K = 7 combinatorics — inheriting the torsion cycles — or only lives on a smooth approximation in which the torsion has washed out into generic continuous holonomy, is exactly the **continuum-limit question**. If the continuum limit preserves the torsion structure, the route closes and "why seven" is answered: seven because the substrate is K = 7. If it washes out the discreteness, the connection's holonomy is generically continuous (Position 4 of §7), no $\mathbb{Z}_7$ arises intrinsically, and the lift's sevenfoldness would have to be donated by the target — leaving the identification weak.

The useful finding is therefore that **"why seven?" reduces to two cross-corpus facts**, neither free-floating: (i) whether the Gate-3 $\mathbb{Z}_7$ is homological torsion or merely a coefficient choice (§6.2), and (ii) whether the continuum limit preserves that torsion in the reverse-map connection's holonomy. The second is the same continuum-limit gap that blocks the decisive precondition (§5) and recurs across the programme. We close neither, but locating "why seven" inside them is progress: a single construction — the continuum limit relating Γ to Γ_vac, together with the torsion reading of the Gate-3 $\mathbb{Z}_7$ — would simultaneously answer why seven, whether the holonomy is quantized, and whether the lift exists.

6.4 Does the corpus favour $\mathbb{Z}_7$-quantized holonomy?

The previous subsections leave the descent open. But there is a separate, weaker argument that can be made without resolving it: even granting that we do not know whether Γ descends from the K = 7 substrate, we can ask which discrete target the corpus would supply *if* it does. The answer shifts the burden of explanation toward quantization, without proving it.

The observation is that the programme contains exactly one persistent source of discrete sevenfold structure: the K = 7 closure algebra underlying the Gate-3 sector. Every known non-trivial residual structure in the closure programme is inherited from this K = 7 organisation; no independent source of discrete valuation has been identified anywhere else in the corpus. Consequently, *conditional on* Γ descending from the same substrate transport processes that generate the Gate-3 closure sector — the very descent §6.2–6.3 leave open — $\mathbb{Z}_7$-valued holonomy is not an arbitrary new condition but the *unique* discrete target already present in the framework. The argument does not supply the descent; it answers a different question (given a descent, why $\mathbb{Z}_7$ rather than some other discrete group?) with: because $\mathbb{Z}_7$ is the only discrete structure the corpus has.

This does not prove quantization. What it does is shift the burden of explanation, by a parsimony argument:

- A theory in which Γ exhibits generic continuous holonomy *while* the substrate simultaneously supports a persistent $\mathbb{Z}_7$ closure sector requires **two** independent residual structures — one continuous (Γ's curvature), one discrete (the Gate-3 sector) — with no stated relation between them.
- A theory in which Γ inherits $\mathbb{Z}_7$ quantization from the same K = 7 closure requires **one** residual structure, appearing in two descriptions.

By structural economy, the second is preferable as the hypothesis to pursue — unless contradicted by further analysis. We are careful about the status of this move: parsimony is a guide to which hypothesis to *pursue*, not evidence that it is *true*. A skeptic may grant the entire argument and still maintain that the substrate simply carries two residual structures; economy is a methodological preference, not a constraint nature is obliged to respect. The argument therefore raises the prior on quantization and places the burden on the two-structure alternative; it does not discharge the question.

Observation 6.1 — The corpus's stance on quantization

The corpus does not force $\mathbb{Z}_7$ quantization of the reverse-map holonomy. But it identifies K = 7 closure as the only presently known source from which such quantization could naturally arise, and by structural economy a single K = 7-inherited residual structure is preferable to two independent ones — conditional on the descent of §6.2 and defeasible under further analysis.

This is the answer to the reader who, told the lift needs $\mathbb{Z}_7$ quantization, asks "why should I expect that?" The expectation is grounded not in the assumption that the lift exists, but in the corpus's track record: every residual discrete structure discovered so far originates from K = 7 closure, so K = 7 is where a reverse-map quantization would be expected to come from too. That is a reason to expect quantization — not a demonstration of it.

6.5 Status

The sevenfoldness on the connection side is therefore **open**, but structured: it is donated-by-target if the continuum limit washes out the discrete structure, and intrinsic-and-shared if the limit preserves it. The identification is a genuine one (two independently sevenfold structures with a common cause) precisely in the case where the continuum limit preserves K = 7 — and an artefact of relabelling otherwise. Which holds is not decided here. What is established is that the question has a definite locus (the continuum limit), that the corpus supplies a positive reason to expect the favourable case (§6.4: K = 7 is the only available source of discrete valuation, and economy favours one structure over two), and that the strength of the whole identification rests on the descent that would convert this expectation into a result.

---

7. The Four Positions — Classifying by Holonomy, Not Curvature

A consequence worth drawing out: the lift question does not inherit the SRP paper's binary readout fork unchanged. The right classifying axis is **holonomy valuation**, not curvature. The SRP fork classified by curvature (Branch A: $\Omega = 0$; Branch B: $\Omega \neq 0$). But §4 showed the lift turns on whether the *holonomy* is $\mathbb{Z}_7$-quantized, which is not the same axis. Crossing curvature (zero / non-zero) with holonomy (quantized / continuous) gives four positions:

1. **$\Omega = 0$ everywhere, trivial holonomy.** The lift is vacuous: $\kappa = 0$. SRP Branch A in the strict sense (fully path-independent); Ω identified with the *trivial* closure class — a lift that exists but carries no structure.
2. **$\Omega = 0$ everywhere, non-trivial $\mathbb{Z}_7$ holonomy** (flat-mod-$\mathbb{Z}_7$, §4.2). The lift exists and is non-trivial: $\kappa \neq 0$ equals Ω's residual holonomy on non-contractible loops. *Path-independent continuously* (no local curvature anywhere) yet *path-dependent discretely* (a sevenfold residue on non-contractible loops). The simplest non-trivial lift.
3. **$\Omega = 0$ on the accessible region, quantized flux confined to holes/torsion cycles** (the Aharonov–Bohm / Dirac-flux pattern). The lift exists: the connection is flat everywhere the loops can actually be deformed, with quantized $\mathbb{Z}_7$ flux threading inaccessible cycles (punctures, holes, or H_1 torsion). The holonomy is deformation-invariant on the accessible loops — homotopic loops enclose the same confined flux — and lands in $\mathbb{Z}_7$. The lift exists for the same reason as Position 2; this differs only in *where* the sevenfold structure sits (threaded through holes rather than carried by a globally flat connection on a multiply-connected accessible space).
4. **$\Omega \neq 0$ on the accessible region** (generically curved). The lift fails (Proposition 4.1, and §4.3): free curvature on the region the loops traverse destroys deformation-invariance — two homotopic loops bounding a surface Σ differ in holonomy by the flux $\int_\Sigma \Omega \neq 0$ — so holonomy is not a function of homotopy class and *no* class κ exists, regardless of the substrate axioms. Ω and κ are distinct (the Distinctness Alternative). SRP Branch B proper.

One point about Position 3 must be made precise here, because it bears on coherence with §4.3. Position 3 must **not** be read as "free non-zero local curvature on the accessible region." That reading would contradict §4.3 outright: §4.3 grounds the existence of κ in the deformation-invariance of holonomy, which holds only where curvature vanishes on the swept region; free curvature there destroys the invariance and leaves no class at all (this is precisely Position 4). The only coherent curved-but-lifting case is the confined-flux pattern stated above — flat on the accessible region, curvature (flux) confined off the deformable cycles. So Position 3 is **nearer-merged with Position 2** than a side-by-side listing suggests: both require flatness on everything the loops can sweep across, and they differ only in whether the sevenfold holonomy comes from the topology of a flat connection on a multiply-connected space (Position 2) or from quantized flux through holes the loops cannot cross (Position 3). The lift exists in Positions 2 and 3 — wherever holonomy is $\mathbb{Z}_7$-quantized *and* deformation-invariant on the accessible loops — and fails in Position 4. Position 1 is the trivial-lift boundary.

The genuinely new structure this analysis surfaces is therefore the *quantized-holonomy* class (Positions 2 and 3 together): connections the SRP binary fork could not name, because that fork classified by curvature alone and quantized-deformation-invariant holonomy cuts across it. Position 2 (flat-mod-$\mathbb{Z}_7$) is the cleanest representative, but Position 3 (confined flux) lifts equally. Crucially, the boundary that *kills* the lift — free curvature on the accessible region — is fixed by flat-connection mathematics, not by the substrate axioms: wherever loops can be freely deformed across curvature, homotopy-invariance fails and no class exists whatever the axioms say.

If the Continuum-Lift Hypothesis is to hold non-trivially, the substrate must realise Position 2 or Position 3 — i.e. the readout holonomy must be $\mathbb{Z}_7$-quantized. That is a sharper claim than "Ω and κ are the same," because it specifies the structural condition (quantized holonomy) exactly.

What would refute the lift, concretely. The lift's observational contact is through reconstruction path-dependence (SRP §9.3): the holonomy $H(\ell)$ is, physically, the discrepancy between committed records reconstructed along two different paths to the same event. The class of measurements that would register it is therefore *path-comparison* measurements — reconstruct a committed record by two routes and compare. Three outcomes map to the positions: (i) reconstructions always agree, across all loops including those that would enclose curvature — full path-independence, Position 1 (vacuous lift); (ii) reconstructions disagree by one of exactly seven discrete amounts, loop-dependent but quantized — sevenfold path-dependence, Positions 2–3 (non-trivial lift); (iii) reconstructions disagree by a continuously-varying amount — continuous holonomy, Position 4 (lift excluded, Ω and κ distinct). Outcome (iii) refutes the lift outright; outcome (ii) confirms a non-trivial lift without yet distinguishing Positions 2 and 3 (that needs a local curvature measurement on the accessible region). We are explicit that this is schematic: realising such a measurement is downstream of the continuum-limit construction that would relate the abstract holonomy to a physical reconstruction procedure, which is unbuilt. The falsifiability is therefore a statement about *what kind* of measurement bears on the question, not yet a specified protocol — but it is a definite kind (path-comparison of reconstructed records), not a gesture.

What would select the Distinctness Alternative. For symmetry — and because a disfavoured horn left thin reads as thumb-on-scale — we state what would establish the *other* horn, that Ω and κ are genuinely distinct. Three findings, any one of which would favour distinctness: (a) the homological computation of §6.2 returns torsion-free $H_1(\Gamma\_vac; \mathbb{Z})$, so the Gate-3 $\mathbb{Z}_7$ is a coefficient choice with no connection-side counterpart; (b) the continuum limit relating Γ to Γ_vac washes out the discrete torsion structure, so even a torsion-bearing complex yields generically continuous readout holonomy; (c) a path-comparison measurement returns outcome (iii) above, continuous path-dependence. Findings (a) and (b) are theoretical and internal to the corpus; (c) is observational. Distinctness is not the residual hypothesis one is left with after failing to prove the lift — it is a positive possibility with its own selecting evidence, and any of these three would establish it as decisively as torsion-plus-quantization would establish the lift.

We do not claim the substrate realises any particular position. We claim that $\mathbb{Z}_7$-quantized holonomy (Position 2 or 3) is the condition under which the hypothesis holds non-trivially, that the corpus neither forces nor excludes it, and that deciding it is the content of the open problem — with the selecting evidence for each horn now stated symmetrically.

---

8. Open Questions

1. **The decisive precondition.** Is the readout holonomy hol(Γ) $\mathbb{Z}_7$-quantized on the admissible loops of $\pi_1(\Gamma\_vac)$? This — not flatness specifically — is what decides the lift

(§4.3, §7): quantized holonomy (Position 2 or 3) gives a canonical κ; continuous holonomy (Position 4) gives none. Flat-mod-$\mathbb{Z}_7$ is the simplest quantized case but not the only one. The question is open.

2. **Continuum limit.** The reduction map (§4.3) presupposes a construction relating the smooth connection Γ to the simplicial complex Γ_vac — the same continuum-limit gap recurring across the programme. Until it is built, the projection H(ℓ) → κ is defined only schematically. Can the limit be constructed, or the lift reformulated directly on the simplicial complex?
3. **Topological support for confined flux.** Positions 2 and 3 both require flatness on the accessible region; they differ only in whether the sevenfold holonomy comes from a flat connection on a multiply-connected space or from quantized flux threading holes/H_1 torsion cycles. Part of this is already settled by flat-connection mathematics, not the axioms: free curvature on the accessible region admits no class (Position 4), full stop. The genuine residual question is narrower — does the substrate possess the topological structure (non-contractible cycles, H_1 torsion) needed to *support* either non-trivial flat holonomy or confined flux at all? If π_1(Γ_vac) and H_1(Γ_vac) are trivial, only Position 1 (vacuous lift) survives; the non-trivial lift requires the vacuum complex to be topologically non-trivial in the right way.
4. **Observable signature.** SRP §9.3 identifies reconstruction path-dependence as the observable distinguishing its branches. Does quantized-holonomy readout leave a *distinct* signature — a discrete, sevenfold path-dependence — separable in principle from both continuous curvature (Position 4) and full path-independence (Position 1)? And can Positions 2 and 3 be distinguished observationally, or only by local curvature measurement?
5. **Direction of the lift.** The hypothesis is framed as Ω → κ (continuous lifts to discrete residue). Is the converse framing — κ as fundamental and Ω as its continuum approximation — equally consistent, and does the corpus prefer one direction? Given the K = 7 substrate is fundamentally discrete, κ-as-primary may be the more natural reading.
6. **Why seven, intrinsically?** Does the reverse-map connection acquire $\mathbb{Z}_7$ holonomy from the K = 7 substrate independently of assuming the lift (§6)? This reduces to *two* cross-corpus facts, both required: (i) that the Gate-3 $\mathbb{Z}_7$ is homological torsion in H_1(Γ_vac; $\mathbb{Z}$), not merely the coefficient group (§6.2) — the single computation Tor(H_1(Γ_vac; $\mathbb{Z}$)) settles this; and (ii) that the continuum limit relating Γ to Γ_vac preserves that torsion in the holonomy (§6.3, the continuum-limit question of Q2). The sevenfoldness is intrinsic and shared — and the identification genuine — only if both hold; if either fails, the $\mathbb{Z}_7$ is donated by the target and the identification is weak.

The contribution of this paper is to have converted "are Ω and κ the same?" into question 1 — a single, sharp, checkable condition ($\mathbb{Z}_7$-quantized holonomy) — and to have established what the corpus fixes around it (the $\mathbb{Z}_7$ target, the reduction machinery, and Proposition 4.1's exclusion of the generically-curved case) and what it leaves open (whether the holonomy is quantized, entangled with the continuum limit). The identification is neither asserted nor refuted; it is made precise.

The single most consequential thing the paper establishes, restated to close: the strength of the whole identification is governed by one computable invariant, Tor(H_1(Γ_vac; $\mathbb{Z}$)) (Central

Finding). If that torsion is trivial, the Gate-3 $\mathbb{Z}_7$ is a coefficient label and the identification is weak; if it equals $\mathbb{Z}_7$, the sevenfoldness is topologically real, the connection can inherit it, and the identification is genuine — the branch in which this line of work becomes interesting. That computation is finite, mechanical, and runnable now on a complex already in hand; it is the decisive next step, and it determines not whether the lift is finished — the continuum limit (Q2) governs that — but whether the lift is real enough to be worth finishing.