Phase as Commitment Memory

Persistent Closure Holonomy as the Physical Origin of Quantum Phase

Keith Taylor — VERSF Theoretical Physics Programme

General Reader Summary

The idea this paper explores

The strange "phase" that quantum theory runs on may not be a basic ingredient of reality at all. It may be the trace left behind by events that can never be undone — reality's faint, permanent memory that something was decided.

Quantum physics is built on a quantity called *phase*. It is what makes particles interfere, what gives quantum behaviour its waviness, and it sits at the centre of nearly every calculation. Yet physics has never said *why* phase exists. It is simply written into the theory at the start, and the theory explains how it behaves but never where it comes from.

This paper asks where it comes from — and offers a candidate answer that runs in an unfamiliar direction.

The VERSF programme is built on the idea that reality is made of *commitments*: events in which something becomes definite and cannot afterward be undone. A commitment is irreversible the way a shattered glass is irreversible — the universe can go forward from it but never back through it. The proposal here is that each such commitment leaves a permanent mark, not a record of *what* happened or *when*, but a faint topological trace that something happened at all. Reality, on this view, carries a kind of memory.

That memory is unlike anything we normally call memory. It stores no events, no order, no details. It is a running tally — and a peculiar one, because it works like a clock face: it counts round in sevens, so that after seven it returns to zero the way a clock returns to twelve. Seven aligned commitments therefore cancel back to nothing, and the tally certifies that commitment occurred but never says how much. (Strictly there is one such clock for each independent way of looping around committed structure, so the full memory is a small handful of these clocks rather than a single one — but the single-clock image is the right one to hold.) It is also not kept in any one place. You cannot point to where it is stored; it shows up only when you travel in a loop and come back to find something has shifted. That "shift around a loop" is exactly the mathematical shape of quantum phase — which is the clue the paper follows.

If the idea holds, the consequence is a reversal of the usual story. We tend to think the past survives because records of it survive — photographs, fossils, memories, notes. This paper suggests the opposite: records survive *because* the past is already fixed into the shape of reality by irreversible commitment. Memory would come first, and records second.

Two honesties are kept throughout. First, the paper does not claim this is proven. It rests on one open question — whether the permanent trace, which exists on one side of the framework, also shows up in the everyday refinement processes the rest of physics is built from — and the paper marks clearly that this single question decides everything. Second, even if it holds, what survives is modest: not a cosmic diary, but a small, cancelling count of net commitment, standing outside time rather than ticking along inside it. What makes that count interesting is not its size but its kind. Every quantity physics currently has describes *what exists right now*. This one would describe *what happened* — and no quantity in established physics is like that. It would be the first: a quantity whose very existence depends on the past being fixed.

Abstract

The origin of phase remains one of the least physically interpreted structures in quantum theory. Standard formulations assign complex amplitudes and phase factors as primitive ingredients, while explanations of interference, holonomy, and gauge transport presuppose phase rather than deriving it.

This paper explores a different possibility emerging from the VERSF programme.

Recent results suggest that irreversible commitment may generate a persistent closure residue represented by nontrivial closure holonomy. The Gate-3 programme identifies this residue as a $\mathbb{Z}_7$-valued **closure transport residue** — a holonomy grading carried by protected non-bounding cycles, inherited from the $K = 7$ simplicial architecture — embeddable into a U(1) phase structure through the seventh roots of unity. (This $\mathbb{Z}_7$ transport residue is a distinct object from the $\mathbb{Z}_2$ fold register settled earlier in the Gate-3 arc; the two are disambiguated in §3 and coexist rather than compete.) Independently, the Fact Momentum and BCB programmes construct a continuum U(1) connection acting on a phase-bearing amplitude and identify closure holonomy with the flat sector of the resulting transport theory — the construction in which the discrete residue reads out as a continuous holonomy.

These observations motivate the **Memory Residue Hypothesis**:

Persistent closure holonomy generated by irreversible commitment is the substrate origin of phase.

The claim is not that quantum mechanics has been derived. Nor is it that phase and memory are identical objects. The proposal is that quantum phase admits a physical interpretation as the continuum manifestation of accumulated commitment residue. Closure holonomy is identified

with the *flat sector* of the transport theory, which covers the geometric phase directly; the account extends to phase as such because, time being emergent in VERSF, no separate fundamental dynamical phase remains — though whether flat-sector holonomy *reproduces* the effective dynamical phase the continuum exhibits is itself an open step, marked as such (§5.3).

We formulate the hypothesis precisely, identify the assumptions on which it depends — in particular the two distinct inputs (irreversibility, native; and essentialness of the discarded region, inherited) that the persistence result requires — and investigate its implications for interference, holonomy, global information, and the Born-rule programme.

The resulting picture suggests that phase may not be primitive. Instead phase may emerge as the transport representation of persistent topological memory generated by irreversible commitment.

Epistemic markers: ***(established)*** *for results inherited from prior VERSF papers;* ***(conditional)*** *for results holding under a stated assumption;* ***(conjectural)*** *for interpretive proposals;* ***(open)*** *for what remains undecided.*

Table of Contents

---

1. Introduction

Quantum theory is built upon phase. The complex phase of the amplitude governs interference, transport, gauge structure, and geometric holonomy. Yet the physical origin of phase remains obscure: phase is normally introduced axiomatically. The formalism explains how phase behaves; it does not explain why phase exists.

The present paper investigates a possible origin. The proposal emerges from a convergence of three independent strands of the VERSF programme:

- irreversible commitment and Fact formation;
- Gate-3 closure holonomy;
- the continuum U(1) transport structure appearing in Fact Momentum.

Individually these programmes address different questions. Together they suggest a common interpretation: persistent commitment residue may appear macroscopically as phase.

The objective of this paper is not to establish that conclusion as a theorem. It is to determine whether the interpretation is mathematically consistent, physically meaningful, and capable of explaining features of phase that otherwise appear primitive — and, equally, to mark plainly the one open question on which the whole interpretation rests.

1.1 Why phase is the remaining primitive

It helps to see why this paper exists at all — why phase, specifically, is the structure worth chasing. The earlier VERSF programmes addressed the foundational ingredients of the quantum side one at a time, and between them they have accounted for most of what quantum theory ordinarily assumes. The ODG and OIP programmes addressed probability — how a probability is extracted from an amplitude. Fact Momentum addressed transport — how amplitudes are carried and compared. Gate-3 addressed persistent topological residue — the protected structure irreversible commitment leaves behind.

What none of these supplied was a physical interpretation for *phase itself*. Phase is present throughout the framework — it is in the amplitude Fact Momentum transports, in the interference ODG and OIP weigh — but it has always entered as inherited mathematical structure rather than as explained physics. Every neighbouring primitive has been given an origin; phase has been used everywhere and grounded nowhere.

The Memory Residue Hypothesis is therefore not introducing a new physical ingredient into the framework. It is the attempt to supply an origin for the *last major primitive still unexplained* on the quantum side. That is the gap this paper addresses, and it is why the question "why does phase exist?" is not idle curiosity but the natural next item on the programme's agenda.

---

2. The Problem of Phase

The role of phase in quantum theory is well known. Observable interference depends not only on amplitudes but on relative phases; gauge theories describe transport through phase rotations; geometric phases appear through holonomy. Nevertheless phase itself is typically treated as fundamental.

The standard question is: *why does phase influence physical outcomes?* The question considered here is different: *why does phase exist at all?*

A satisfactory answer should explain four things:

1. why phase possesses a transport character;
2. why phase accumulates around loops;
3. why phase naturally participates in interference;
4. why phase is associated with global rather than purely local structure.

These are precisely the properties exhibited by closure holonomy. The remainder of the paper asks whether that coincidence is structural rather than accidental.

---

3. Closure Residue and Persistent Holonomy

The Gate-3 programme identifies a closure residue associated with nontrivial cycles. The residue is represented by closure holonomy rather than by local state variables, and two of its properties matter here.

A naming point must be settled before anything else, because the programme contains two distinct $\mathbb{Z}$-graded closure objects and conflating them would put this paper in apparent contradiction with settled work. Earlier in the Gate-3 arc, the *register of the closure charge* was settled as $\mathbb{Z}_2$ — a committed bit, fixed by fold uniqueness and the saturation theorems. That $\mathbb{Z}_2$ object is the **fold register**, and this paper does not touch it. The object this paper is about is different: a **$\mathbb{Z}_7$-valued closure transport residue** (equivalently *K = 7 transport residue* or *closure transport holonomy*), a holonomy grading on the vacuum transport sector that descends from the K = 7 simplicial architecture (§5.1). The two are not rivals and do not compete: the $\mathbb{Z}_2$ fold register records the committed bit of a fold; the $\mathbb{Z}_7$ transport residue grades the holonomy of transport around protected cycles. They live on different structures (fold register vs transport sector), answer different questions (which bit was committed vs what holonomy a loop carries), and coexist without tension. Throughout this paper, "closure residue," "closure holonomy," and "transport residue" refer to the $\mathbb{Z}_7$ transport object; where the $\mathbb{Z}_2$ fold register is meant it is named explicitly. The phrase "closure charge" unqualified is avoided precisely because it already denotes the $\mathbb{Z}_2$ register in the home papers.

With the naming fixed, two properties of the $\mathbb{Z}_7$ transport residue matter here.

First, the residue is topological. It is associated with loops rather than points: its value is read around a closed cycle, and it vanishes on any loop that can be contracted.

Second, the residue is persistent — but persistence rests on two inputs, not one, and the distinction is essential to state honestly. Under the persistence branch of the Gate-3 analysis, a primitive-Fact cycle γ_D is permanently non-bounding *on the closure side*. The reason is not irreversibility alone. A committed Fact discards a region D, and a loop drawn around D can fail to contract for two separate reasons that must both hold:

- **Irreversibility** forbids contracting the loop *through* D — a contraction sweeping across the discarded region would reach back into it and undo the commitment, which is the one motion the substrate cannot perform. This is the native part: it follows from commitment being irreversible.
- **Essentialness** forbids contracting the loop *around* D on the far side. Discarding a region does not by itself trap a loop: on a closed surface a small discarded patch leaves its boundary free to slide off the other way and shrink to nothing. The loop is trapped only if the discarded region is *genuinely encircled* — if there is no way around it either. That the discarded region is essential in this sense is an inherited input from the closure papers, and it carries whatever standing those papers give it.

Only with **both** — irreversibility (no path through) and essentialness (no path around) — is γ_D permanently non-bounding. Stating it as a consequence of irreversibility alone would claim more than the framework proves; the topological weight sits in essentialness, the permanence in irreversibility, and the honest claim needs the pair.

With that pair in hand, any nonzero closure holonomy supported on such a cycle survives — on the closure side. The resulting object has exactly the characteristics expected of a memory residue: it records that commitment occurred, it records no chronology, it is global rather than local, and it survives because commitment cannot be undone and the region it carved out cannot be circumvented. Whether this closure-side permanence also becomes a feature of refinement motion is a separate question, deferred to §10; §§4–9 read the closure-side object as physics under the supposition that it survives.

---

4. The Memory Residue Hypothesis

The central proposal of this paper is the following.

Memory Residue Hypothesis (conditional)

Suppose the persistence branch of the Gate-3 programme is correct — that is, suppose the two inputs of §3 (irreversibility and essentialness) hold and a primitive-Fact cycle survives as a protected, charge-carrying loop. Then irreversible commitment generates a surviving closure residue represented by nontrivial holonomy on protected cycles, and this residue acts as a form of topological memory.

The hypothesis and its consequences are unpacked across §§5–12, but the claim itself can be stated in one place, with its premises made explicit, so that there is a single sentence to point to rather than a thesis distributed through the paper. It is not a mathematical theorem — it proves no new mathematics — but an *interpretation theorem*: a statement of what physical reading the established and conditional inputs license, and under exactly which conditions.

Interpretation Theorem (conditional)

If

1. **primitive-Fact persistence holds** — a primitive-Fact cycle is permanently non-bounding on the closure side, which by §3 requires both irreversibility (native) and essentialness of the discarded region (inherited);
2. **closure holonomy survives refinement** — the surviving cycle is refinement-realised, so the closure-side residue is also a residue of admissible refinement motion (the open question of §10);
3. **the Fact Momentum U(1) transport construction is valid** — the continuum U(1) connection acting on the phase-bearing amplitude is well-founded, so the $\mathbb{Z}_7$ closure transport residue reads out as a continuous U(1) holonomy (§5);

then quantum phase admits an interpretation as the transport representation of accumulated commitment residue.

The three premises are of three different kinds, and keeping them distinct is the point of stating the theorem at all. Premise (1) is part-native, part-inherited; premise (2) is the single open hinge on which everything turns; premise (3) is an inherited construction from a companion programme. The conclusion claims an *interpretation*, not a derivation: it does not assert that phase has been built from commitment, only that — if the three inputs hold — phase *may be read as* the continuum face of persistent commitment memory. Everything that follows is the content of that conditional, and §10 isolates premise (2) as the one input still genuinely undecided.

A word on calling this a *theorem*, since "X admits an interpretation as Y" is not on its face a truth-evaluable conclusion. The theorem-content is not the interpretation itself but the **sufficiency claim**: that premises (1)–(3) are *jointly sufficient* for the interpretation, with nothing further required — no fourth input, no additional bridge. That is a definite, checkable assertion about a minimal premise set, and it is what the label is meant to mark. The reader who grants the three premises is committed to the reading; the reader who rejects the reading must reject one of the three. Stated that way the theorem earns its name without dressing a conditional in mathematical clothing — it is a claim about *what suffices*, and it would be false if some premise beyond the three were needed.

The residue does not preserve chronology. It does not preserve a record of individual events. It preserves only the *net* topological consequence of commitment. The residue therefore remembers *that* commitment occurred without remembering the narrative details of commitment.

This distinction is essential. The proposal is not that reality stores a universal archive; it is that reality retains a persistent topological tally of commitment. The residue is historical in origin but topological in form.

The tally is quantised, and the quantisation has a consequence worth drawing out, because it is the framework's most distinctive and most easily misread feature. The closure transport residue is valued in $\mathbb{Z}_7$ — it counts in sevens, in the clock-arithmetic sense: after seven it returns to zero, the way a clock returns to twelve. Aligned commitments can therefore cancel: seven of them sum back to the identity. The residue is consequently a **one-way witness**. A nonzero tally certifies that genuine commitment has occurred; a zero tally certifies nothing, since it may mean no commitment or it may mean a multiple of seven aligned commitments that have summed back to nothing.

One refinement of the "single tally" picture is needed for precision, because it bears on how much the residue can hold. The charge is not, strictly, one global number. It is read around loops, and there is in general a $\mathbb{Z}_7$ value *per independent protected cycle* — formally an element of $\mathrm{Hom}(H_1, \mathbb{Z}_7)$, one seven-valued reading for each independent way of encircling committed structure. Where there is a single protected cycle this collapses to one running tally, which is the image to keep in mind; in general the residue is a small bundle of such tallies, one per independent loop. Either way the budget is bounded and tiny: a few bits per cycle, scaling with the rank of the relevant homology rather than with the number of events. The substrate keeps a cancelling count of net commitment — not a count of events, and beneath the net no finer record at all.

(The cancellation picture — that *seven aligned* commitments sum to the identity — presupposes that each primitive commitment contributes one unit of the $\mathbb{Z}_7$ charge, i.e. one generator. This unit-charge-per-commitment is an inherited Gate-3 input, not established here; where it fails, the cancellation arithmetic changes accordingly, though the qualitative one-way-witness character survives as long as the charge group is finite.)

A natural follow-on question is how much memory the substrate holds *in total*, since the per-cycle budget is bounded but the total is (bits per cycle) × (number of independent protected cycles) = a few bits × rank H_1. The witness-not-archive framing is secure per cycle regardless; whether it remains apt globally depends on whether rank H_1 is itself bounded in the substrate. Whether the closure papers bound it is an inherited question we do not settle here: if rank H_1 is bounded, the total budget is bounded and the framing holds at every scale; if it grows without limit, the residue could carry unboundedly many independent few-bit tallies, and "tiny" would describe each tally but not their sum. We flag this as **(open, inherited)** rather than assert a bound — the per-cycle claim stands either way, and only the global aggregate turns on it.

---

5. From $\mathbb{Z}_7$ Transport Residue to U(1) Phase

The closure residue appears initially as a discrete object. The Gate-3 construction associates a transport residue with protected cycles, and the resulting quantity is naturally represented by elements of a finite cyclic group: the residue lives in $\mathbb{Z}_7$.

5.1 Why the memory counts in sevens

Before earning the step to a continuous phase, it is worth saying why the charge group is $\mathbb{Z}_7$ and not some other finite group — because a seven that arrived by convenience would make the whole interpretation look arbitrary, and the seven does not arrive by convenience.

The appearance of $\mathbb{Z}_7$ is **not introduced by the present paper** — but it is worth being exact about which part is inherited and which part this section actually contributes, because they are not the same and the genuinely novel step should not wear inherited clothing.

What is **inherited** is the **K = 7 closure architecture**: the substrate's admissible closure geometry possesses a sevenfold structure, fixed by the no-go analysis of the relational substrate elsewhere in the programme. That seven is established and nothing here re-derives it.

What does **not** follow from "K = 7" alone is that the holonomy of the transport sector is graded by the *cyclic* group $\mathbb{Z}_7$. Seven elements generically carry the full symmetric-group action S_7; getting from "seven-fold" to "$\mathbb{Z}_7$" requires the sevenfold structure to be specifically *rotational* — a cyclic ordering closed under a single generating rotation — so that its holonomy classes are clock arithmetic mod 7 rather than arbitrary permutations. It is that cyclic reading of the K = 7 architecture that yields $\mathbb{Z}_7$ and licenses the seventh-roots embedding of §5.2. We flag this as the **interpretive bridge the present section contributes** (or, if the cyclic-holonomy structure of the closure transport sector is already established in a home paper, an inherited input to be cited there — we mark it conditional pending that confirmation, rather than assert it as settled). Either way it is the load-bearing link, and it is a claim about *cyclic* structure, not merely about the number seven.

Granting the cyclic reading, the memory residue counts in sevens not because seven was chosen as a convenient finite group, but because the substrate's admissible closure geometry is sevenfold *and rotational*. The mod-7 arithmetic of commitment memory is then a reflection of the same K = 7 architecture from which the closure transport sector emerges.

This enriches the chain the rest of the section builds. The phase structure is not merely

memory → $\mathbb{Z}_7$ → U(1),

with the seven appearing from nowhere, but

K = 7 closure geometry → $\mathbb{Z}_7$ transport residue → U(1) phase, in which the discrete group is the fingerprint of a specific substrate geometry rather than a free parameter. The seven in the memory is the seven in the closure architecture, seen one layer up.

There is a stronger interpretive observation available here, which we state as such rather than as a result. If the programme is right, the reason quantum phase appears *periodic* may ultimately be that the underlying closure architecture is periodic. On this reading the circle of phase is not fundamental; it is the continuum shadow of the discrete sevenfold closure structure — the U(1) we measure being what the $K = 7$ substrate geometry looks like after emergent time and the continuum limit have smoothed its seven points into a circle. We do not claim this as established; as an interpretive observation it says that periodicity of phase and periodicity of closure would be one fact wearing two descriptions, the discrete one primary.

This is worth stating at full strength, because it is more than the bare inclusion $\mathbb{Z}_7 \subset U(1)$. The inclusion says only that the seven values *fit inside* the circle. The claim here is causal in direction:

phase periodicity $\Leftarrow$ closure periodicity.

The phase circle is periodic *because* closure transport is periodic — the periodicity is inherited, not coincidental. This is the difference between observing that seven points happen to lie on a circle and explaining why there is a circle at all: a primitive U(1) phase has no account of *why* phase should be periodic (it simply is, by postulate), whereas a phase descended from $K = 7$ closure transport is periodic for a reason — it inherits the periodicity of the structure it is the shadow of. The proposal therefore does not merely place the seven inside the circle; it derives the circle's defining feature, periodicity, from the substrate. That is the sense in which "why seven" and "why periodic at all" are the same question with one answer.

5.2 From $\mathbb{Z}_7$ to U(1)

The step from a seven-valued tally to a continuous phase is the load-bearing move of the whole proposal, and it should be earned rather than asserted. A residue that *counts in sevens* is not, on its face, the same kind of object as a continuous phase governing interference; the embedding has to be exhibited, and its continuum half flagged as inherited.

The embedding has two parts.

The discrete half (elementary). The seven values of $\mathbb{Z}_7$ sit inside the phase circle U(1) as the seventh roots of unity,

$k \in \mathbb{Z}_7 \mapsto \exp(2\pi i k / 7) \in U(1), \; k = 0, 1, \ldots, 6.$

This is a faithful group homomorphism $\mathbb{Z}_7 \hookrightarrow U(1)$: the additive structure of the tally (commitments adding, and seven of them cancelling) maps exactly onto multiplication of phase factors (the roots multiplying, and the seventh power returning to 1). The mod-7 cancellation of §4 *is* the closing of the phase circle after seven steps — and, by §5.1, the closing of the closure geometry after its seven rotational steps. So the discrete charge already lives in the phase circle, as a privileged seven-point subset of it.

The continuum half (inherited from Fact Momentum and BCB). Embedding the values into U(1) is not yet the same as a continuous phase that transports and accumulates. That further step — by which the discrete closure transport residue reads out as a *continuous* holonomy of a genuine connection — is supplied by the Fact Momentum programme, which constructs a continuum U(1) connection acting on a phase-bearing amplitude and identifies closure holonomy with the flat sector of that transport theory. We do not re-derive that construction here; we take it as an established input and use only its output: that the seventh-root-of-unity residue of a closure cycle is the discrete shadow of a continuous U(1) holonomy in the continuum description.

There is, moreover, an independent reason the connection must exist and must be U(1), which it is worth importing because it makes the appearance of U(1) non-arbitrary from a second direction. The BCB results establish that amplitude space carries a *single global phase redundancy* — one direction of indistinguishability, along which states differing only by overall phase are physically identical — and that this redundancy admits a unique local gauge-lift. The argument is short: there must exist at least one U(1) gauge field because the single global phase redundancy is the unique direction of indistinguishability in amplitude space, and local comparison of phase between separated points requires that redundancy to be transported — which is exactly a U(1) connection. So U(1) is forced twice over: abstractly, as the unique compact connected one-dimensional completion of periodic transport (Proposition 5.2 below); and physically, as the gauge-lift the unique amplitude-space phase redundancy requires for local comparison (BCB). The Memory Residue Hypothesis does not have to explain *why U(1) exists* — BCB and Fact Momentum supply that. Its role is narrower and complementary: to propose the physical *origin of the phase* whose transport that U(1) describes. One programme explains the transport structure; the other explains the source of the transported quantity.

The relation between the two halves is best read as **charge sourcing a connection**, not as a discrete object somehow becoming continuous. The quantised $\mathbb{Z}_7$ transport residue sources a continuous holonomy the way a quantised electric charge sources a continuous field: the source is discrete, the field it lights up is not, and no category error is involved in their coexistence. This is the precise sense in which "the memory creates the phase" is defensible without overreach. Commitment-memory supplies the *source* — the discrete, commitment-derived $\mathbb{Z}_7$ transport residue; Fact Momentum and BCB supply the *connection* along which that source propagates as continuous phase. The continuum is therefore not claimed to be commitment-derived; what is commitment-derived is the residue that lights it up. Memory provides the seven-point skeleton; the inherited connection provides the circle.

One question the inheritance from Fact Momentum leaves implicit deserves to be made explicit: granted that the continuum readout is *some* continuous phase, why is it U(1) — the circle — rather than some other continuous group? The answer is not free choice; it is forced once the structural requirements on a continuum completion of periodic transport are written down.

Proposition 5.2 (U(1) as the forced continuum completion) (proven, given the stated structural hypotheses)

Suppose the continuum readout of closure transport is required to be (i) a connected Lie group — transport composes and varies continuously; (ii) compact — the transport is periodic, closing on itself rather than running off to infinity; and (iii) one-dimensional / abelian — a *single phase degree of freedom*, i.e. one connection with one fiber and one gauge group, the values commuting as the $\mathbb{Z}_7$ residue does. Then the completion is **U(1)**, uniquely: U(1) is the only compact connected one-dimensional Lie group, and it is the unique such group containing every finite cyclic group $\mathbb{Z}_n$ as a subgroup.

A clarification of (iii), because §4 established a bundle of tallies, not a single value. §4 is explicit that the residue is Hom(H_1, $\mathbb{Z}_7$) — one $\mathbb{Z}_7$ value per independent protected cycle, a bundle scaling with rank H_1. This does **not** conflict with dimension-one, and the distinction is the standard one between a gauge group and its holonomies. Dimension-one is a fact about the *fiber*: there is one U(1) connection, one gauge group, one phase degree of freedom. The many tallies of §4 are not many gauge groups (which would give a torus U(1)^(rank H_1)); they are the many *values* that single connection takes when evaluated on different independent loops. This is exactly electromagnetism, whose gauge group is U(1) — one fiber — while a space of complicated topology carries many independent Aharonov–Bohm phases, one per independent loop. Hypothesis (iii) asserts the single fiber, not a single holonomy value; §4's bundle is the loop-spectrum of that one fiber. So "one connection, many loop-values" reconciles the two: the gauge group is U(1) (not a torus) and the residue is still Hom(H_1, $\mathbb{Z}_7$) (not a single number).

Argument. The classification of one-dimensional connected Lie groups gives exactly two: the line $\mathbb{R}$ and the circle U(1). Hypothesis (ii), compactness — which is the continuum face of *periodicity*, the closure transport returning to its start — excludes $\mathbb{R}$, which is non-compact and models non-periodic accumulation. That leaves U(1). It is abelian (hypothesis iii is then automatic in dimension one) and contains each $\mathbb{Z}_n$ as the nth roots of unity, so it accommodates the $\mathbb{Z}_7$ residue as §5.2 requires. No other group meets all three conditions: drop compactness and $\mathbb{R}$ returns (non-periodic phase); raise the *fiber* dimension and one gets tori or non-abelian groups (multiple or non-commuting phase degrees of freedom), excluded by (iii) — note this is fiber dimension, not the number of independent loops, which is unrestricted and gives the §4 bundle.

The force of the proposition is that the three hypotheses are not stipulations about U(1) but readings of the substrate structure already in hand: connectedness from continuous transport, **compactness from the periodicity §5.1 traced to K = 7**, dimension-one from there being a single closure transport residue. Given those, U(1) is not chosen — it is the only completion available. So the chain

K = 7 closure geometry → $\mathbb{Z}_7$ transport residue → U(1) phase

has every arrow motivated: the first by the closure architecture (§5.1), the second by the embedding (§5.2), and the third by Proposition 5.2 — the circle is forced by the periodicity the seven already carries. (The hypotheses are inherited structural facts about closure transport, not proved here; the proposition shows only that *given* them U(1) is forced, which is where the freedom in the continuum step actually lies.)

With both halves, the residue may be represented as a phase-valued transport quantity, and the chain reads, in its fullest form,

commitment → closure residue → $\mathbb{Z}_7$ transport residue → phase → U(1) transport.

Every arrow now has a rationale attached. The first is the persistence result of §3 (conditional). The second is the K = 7 → $\mathbb{Z}_7$ grading of §5.1, with its cyclic-holonomy bridge owned there. The third is the seventh-roots embedding of §5.2 — the discrete residue *is* a phase, a seven-point subset of the circle. The fourth separates the transported quantity (phase) from the structure that transports it (the U(1) connection): that connection is forced abstractly by Proposition 5.2 and physically by the BCB phase-redundancy argument above. The division of labour is exact — the Memory Residue Hypothesis supplies the first three arrows (the source of the phase); Fact Momentum and BCB supply the last (the transport the phase is carried by).

This five-arrow form is the canonical chain; later sections (§5.3, §13, §14) abbreviate it to three or four arrows for emphasis, collapsing the K = 7 → $\mathbb{Z}_7$ steps or the phase → U(1)-transport steps as the local point requires. These are the same chain at different resolutions, not different chains.

5.3 Flat-sector holonomy and the absence of a separate dynamical phase

A careful reader will object at this point that closure holonomy is identified with the *flat sector* of the transport theory, and that flat-sector holonomy is not all of phase. In standard quantum theory phase comes in two kinds: the *geometric* phase, which is exactly a flat-connection (holonomy) effect of the Aharonov–Bohm and Berry type, and the *dynamical* phase, which is accumulated action — energy integrated against time. Closure holonomy plainly accounts for the first. It does not obviously account for the second, and if dynamical phase were a separate fundamental ingredient the proposal would explain only part of phase, not phase as such.

In most frameworks this objection would force a retreat to "the geometric sector of phase." In VERSF it does not, and the reason is structural rather than convenient. **Time is emergent in VERSF — it is not a substrate primitive.** But dynamical phase is *defined* as energy × time; it presupposes a fundamental time variable to integrate against. A substrate that has no fundamental time has no fundamental dynamical phase to account for separately. What the continuum calls dynamical phase must therefore be an *effective* description that arises together with emergent time, not a second primitive standing alongside the geometric phase.

The consequence is exactly the one the proposal needs. At the substrate level there is only one kind of phase — geometric, holonomic, route-dependent — and the apparent split into geometric and dynamical phase is a feature of the emergent, continuum description, not of the substrate. Flat-sector holonomy is not a *part* of substrate phase to be supplemented later; it is, on this reading, the whole of it. The geometric/dynamical distinction would then be recovered downstream, as emergent time makes an "action accumulated over time" description available.

It is important to mark precisely which half of this is secure and which is not, because it is easy to put the hedge in the wrong place. Emergent time is **established** in VERSF, not assumed here — so the negative half of the argument, *that no separate fundamental dynamical phase remains to be accounted for*, follows directly and is not the conditional part. What is **not** yet shown is the positive, reproduction half: that the single holonomic substrate quantity, read through emergent time, actually *reproduces* what the continuum splits off as dynamical phase. Saying "the underlying quantity was holonomic all along" asserts that reproduction; it does not demonstrate it. The conditional marker therefore belongs here, on the reproduction claim, not on emergent time.

Stated as a target: what a future result must establish is that flat-sector holonomy *generates the correct effective dynamical phase* in the continuum limit — not merely that no separate primitive survives. The distinction matters because the two are not the same. It is conceivable that the effective dynamical phase carries interference content that is *not* recoverable from the flat sector — that emergent time introduces phase structure the substrate holonomy alone cannot reproduce. If so, the gap the emergent-time argument closes at the level of *primitives* would reopen one level down, at the level of *effective content*. The emergent-time move removes the easy version of the objection (a second fundamental phase); the reproduction claim is the hard version, and it remains **(open)**. We mark it so rather than letting "holonomic all along" stand as if proven.

The chain just displayed — commitment → closure residue → closure holonomy → U(1) phase — transforms the interpretation of phase. Conventionally phase is introduced as a primitive mathematical object. The memory-residue picture instead suggests that phase inherits its existence from commitment:

The phase does not create the memory. The memory creates the phase.

This is the first sense in which phase may be interpreted physically rather than axiomatically. We mark the full identification **conjectural**; what is on firmer ground is the structural claim that the surviving quantity has the formal character of a phase — periodicity, holonomy, transport-dependence — and that its discrete signature embeds in the phase circle exactly as the cancellation rule requires.

---

6. Transport and Interference

The defining property of phase is transport. Phase is not measured at a point; it is measured comparatively. What matters is the accumulated phase difference between alternative routes.

Closure holonomy possesses exactly this character. The quantity is not local; it accumulates through transport; it is detected through loops. The similarity is structural, not metaphorical: a holonomy *is* a comparison of transport along different routes between the same endpoints.

If phase arises from persistent closure residue, then interference acquires a natural interpretation — and the interpretation must be stated carefully, because the obvious reading is the wrong one. Interference does **not** arise because multiple histories are stored and then interact; the memory-residue hypothesis stores no chronology, so there are no stored individual histories to interact. Instead, two transport paths between the same endpoints accumulate the residue differently, and their difference is a closed-loop holonomy — a single accumulated charge read around the loop the two paths enclose. Interference becomes sensitivity to that accumulated commitment residue: the path-dependence of one charge, not the interaction of many records.

The interpretation remains speculative. But it replaces an unexplained primitive structure with a physically motivated one, and it does so without smuggling back the archive the hypothesis explicitly denies.

This is the point to settle the account opened in §2, where phase was given four properties an origin ought to explain. Closure holonomy supplies all four, and the correspondence is exact rather than suggestive:

1. *Transport character* — the transport residue is read by transporting around a loop, not measured at a point; holonomy is transport by definition.
2. *Loop accumulation* — the charge accrues around closed cycles and vanishes on contractible ones, which is precisely how phase accumulates around loops.
3. *Interference participation* — interference is the route-difference of accumulated charge (above), so the residue participates in interference in exactly the way phase does.
4. *Global association* — the charge lives in $\mathrm{Hom}(H_1, \mathbb{Z}_7)$, a global topological invariant attached to loops rather than points, matching phase's global character.

The four properties that made phase look like a primitive with no account are, one for one, the defining properties of closure holonomy. That is the structural case for reading the second as the origin of the first.

6.1 The signature of phase, and what carries it

The correspondence above can be sharpened from an observation into a characterisation. The four properties are not an arbitrary list; they are jointly the *structural signature* of phase — the features any substrate object must have if it is to be a candidate origin of phase at all. This lets the claim be upgraded from "closure holonomy resembles phase" to "closure holonomy is the only thing in the programme that does."

Before characterising the signature, it is worth posing the explanatory problem in its most direct form, because that is what the propositions of this section answer.

Proposition 6.0 (Phase Needs a Carrier) (framing; the disjunction is exhaustive, the candidate count relative to identified objects)

Any account of phase must identify *something that carries it*. The options are exhaustive: either (a) phase is primitive — carried by nothing, taken as given; or (b) phase is carried by topology

— a holonomy of some connection; or (c) phase is carried by some other substrate structure. There is no fourth option, because to explain phase *is* to say what bears it, and a quantity is either unborne (a), borne by topology (b), or borne by something non-topological (c).

Why this reframes the problem. The standard treatment is silent on the choice — it writes phase into the formalism and never asks what carries it, which is option (a) by default rather than by argument. Once the question "what carries phase?" is asked explicitly, (a) is revealed as a refusal to answer it, and the live explanatory options are (b) and (c). The work of §6.1 is then to show that among *identified* substrate structures, exactly one — closure holonomy — is a working instance of (b), and none is an instance of (c). The question the proposal puts to the reader is therefore not "is this account correct?" but the sharper *if not closure holonomy, what carries phase?* — to which the only present answers are "nothing (it is primitive)" or "something not yet identified." This attacks the explanatory problem head-on, where Proposition 6.1 attacks its structural form.

The next proposition establishes that the carrier's *gauge-invariant content* cannot be local — which already forces (b) or a non-local version of (c), and rules out the most natural form a rival carrier might take. The phrasing matters here, and is sharpened below to avoid an apparent collision with §5.2.

Proposition 6.0b (No Local Memory) (proven)

No gauge-invariant local observable can carry commitment memory. Any quantity that records the commitment residue gauge-invariantly must be global.

A clarification first, because §5.2 appears to say the opposite. §5.2 has the residue propagate along a U(1) *connection*, and a connection is local data — a 1-form, specified pointwise (up to gauge). So the carrier is, at one level, perfectly local: the connection lives at points. The claim of this proposition is not that no local object is involved — the connection plainly is — but that no *gauge-invariant local observable* recovers the residue. The connection carries the memory locally but gauge-dependently; it records nothing gauge-invariant at any single point, because a gauge transformation can change its value there freely. The gauge-invariant content — the thing that is actually the residue, the actually-measurable memory — appears only when the connection is integrated around a closed loop, and that is irreducibly global. This is exactly the situation in electromagnetism: the vector potential is local but gauge-dependent; the Aharonov–Bohm phase, the gauge-invariant observable, is the loop integral and is global. Local data underlies the holonomy; gauge-invariant local data does not exist for it.

Argument. A gauge-invariant local observable would assign the residue a value at a point (or in a region) unchanged by gauge transformations. But the commitment residue has a property no such observable can have: it survives topological deformation, is read around loops, and depends only on the loop's class — not on any local detail, and not on gauge. Any quantity defined pointwise and gauge-invariantly is, by construction, insensitive to the global loop class that the residue is precisely sensitive to; conversely the connection, which *is* sensitive locally, is gauge-dependent and so carries no gauge-invariant value at any point. So no gauge-invariant local observable can equal the residue: pointwise-gauge-invariant and loop-class-sensitive are incompatible. The residue's gauge-invariant content must therefore be global — a holonomy. This is *why* commitment memory appears as holonomy specifically: holonomy is the canonical object that is carried by local (gauge-dependent) connection data yet is itself global and gauge-invariant, which is exactly the residue's profile.

This is more than a restatement of "the residue is global." It is a *reason* the residue's observable content must be global: not an observed feature but a forced one, following from the incompatibility of gauge-invariance-at-a-point with sensitivity-to-loop-class. Commitment memory could not have been a gauge-invariant local field even in principle — the holonomic form is the only kind of object that can be locally sourced yet globally and gauge-invariantly read. And this reconciles §6.0b with §5.2 rather than opposing it: §5.2 supplies the local (gauge-dependent) connection that carries the residue, §6.0b explains why its *measurable* content is nonetheless forced to be the global loop integral. Local carrier, global observable — the two sections describe the two halves of one holonomy.

Proposition 6.1 (Phase Characterisation) (proven, relative to identified candidates)

Let a quantity be a *candidate origin of phase* only if it (1) is transported rather than locally measured, (2) accumulates around loops, (3) participates in interference through route-dependence, and (4) is associated with global topology rather than local state. Then any substrate object that can serve as the origin of phase must possess all four properties — and among the objects so far identified in the VERSF programme, closure holonomy is the only one that does.

A word on "proven" here, in the same spirit §4 defended "theorem," since the rigor-signal should stay consistent across the paper. The proposition has two parts of different character. The *necessity* of the four properties is **proven in the strict sense** — it is definitional: the four unpack what "phase" operationally means, so a quantity lacking one cannot be phase's origin, with no inductive gap. The *uniqueness* claim is **proven only in the survey sense** — a complete inventory check over the programme's *identified* substrate quantities, true of that list and explicitly silent about anything not yet on it. So "proven, relative to identified candidates" means precisely: necessity proven outright, uniqueness proven over the present inventory. That is what the qualifier marks, and it is why the result is strong without being an absolute uniqueness theorem. (The same reading attaches to Proposition 6.2's "proven, relative to identified sources": its disjunction is exhaustive by strict logic, while its bite depends on the same inventory survey.)

Argument. The four conditions are not extra requirements imposed on a candidate; they are what phase *is*, operationally. A quantity failing (1) is a local observable, and a local observable cannot reproduce a quantity measured only comparatively; failing (2) it does not accumulate as phase accumulates; failing (3) it cannot drive interference; failing (4) it is a local-state quantity and cannot match phase's global character. So the four are necessary, not optional: anything lacking even one is disqualified as an origin of phase. Closure holonomy possesses all four (§6, items 1–4). A survey of the other identified substrate quantities — local refinement data, capacity, arity, ordering, and the rest of the eliminated candidates of the Reversible-Connectedness programme — finds none with the full signature: each fails (4), being local, or fails (1)–(2), not being a transported loop quantity. Closure holonomy is therefore the unique *identified* candidate.

The proposition does **not** prove the identification of phase with closure holonomy, and the qualifier "relative to identified candidates" is essential and load-bearing. It is not a uniqueness theorem in the absolute sense — it does not show that no as-yet-unidentified object could carry the signature. What it shows is that, of everything the programme has so far put on the table, exactly one object has the structure phase requires. That upgrades the paper's claim from

closure holonomy looks like phase

to

closure holonomy is the only thing we have found that looks like phase —

which is the difference between an interesting analogy and a genuine explanatory candidate.

Proposition 6.2 (No Free Phase) (proven, relative to identified sources)

If persistent closure residue does not exist — if the open question of §10 resolves negatively — then the programme possesses no presently identified substrate source for phase. Consequently one of three things must hold: (a) phase is primitive, taken as an unexplained ingredient; (b) persistent closure residue survives, and is the source; or (c) some substrate source not yet identified supplies phase.

Argument. By Proposition 6.1, closure holonomy is the only identified candidate with the phase signature. If the residue does not survive refinement (§10), that candidate is unavailable, and no other identified object has the signature. So phase must then either be accepted without a substrate origin (a), or await a source the programme has not yet found (c) — there is no third *identified* option, which leaves (b) as the only currently-available positive account. The trichotomy is exhaustive because (c) is the explicit catch-all for the unknown.

The force of this is to remove the reading on which the hypothesis is merely an optional curiosity. A reader inclined to think "perhaps memory residue is interesting but unnecessary" is confronted with the fork: *if not this, then what?* The honest answer includes the open (c) — the programme cannot claim to have ruled out every undiscovered alternative — but the point stands that, as matters actually stand, the choice is between treating phase as primitive and taking the residue as its origin. The hypothesis is not one account among several on offer; it is the only substrate account on offer. That is a sharper reason to care whether it survives than any of its individual consequences.

A familiar anchor makes both propositions less exotic than they may sound, and it is worth pressing it past analogy into something closer to evidence. Geometric (Berry) phase already demonstrates, within entirely standard and experimentally confirmed physics, that phase *can* arise from topology rather than from a primitive postulate — a holonomy of a connection, read around a loop, with no local source. This is not a VERSF claim; it is textbook quantum mechanics, measured in countless experiments.

The consequence for the present proposal is a consistency check it passes rather than a resemblance it invokes — but the size of the remaining step must be stated honestly, because it is larger than "scope" alone. The statement "phase can originate from topology" is, in principle, *already established* — physics contains worked examples. So the Memory Residue Hypothesis does not have to defend the *mechanism* — topology yielding phase values — as a novel or speculative possibility; that mechanism is real and observed.

But Berry phase establishes less than the hypothesis needs, and the gap should be named precisely. A Berry phase is a topological contribution to particular phase *values*, computed *within* a quantum theory whose U(1) phase manifold is already primitive. It demonstrates topological sources of values *inside* the circle; it presupposes the circle. The hypothesis claims something categorically larger: that the circle *itself* — the U(1) phase manifold within which Berry phases are computed — is emergent from K = 7 closure. So the extension is not only along a scope axis ("more phase values are topological"); it is a change of level: from *values within U(1)* to *U(1) itself*. Berry exhibits topology producing phase against a background phase structure it assumes; the hypothesis asks for topology producing the background phase structure.

What Berry phase therefore supplies is genuine but bounded: it is a *consistency anchor*, proof that topology-yielding-phase is a real and observed mechanism, so the hypothesis is not inventing a new kind of physics at the level of values. What it does **not** supply is the harder half — it cannot witness the emergence of U(1) itself, because it presupposes U(1). The honest statement is: the mechanism (topology → phase values) is established; the extension (topology → the phase manifold) is the speculative content, and it is more than scope. The slogan "phase is topological all the way down" names exactly this larger claim — *all the way down* meaning down past the values to the circle they live on — and it is the right slogan precisely because it marks that the hypothesis reaches a level Berry phase does not.

This sharply narrows what is actually being proposed without overstating how small the remaining step is. A reader who accepts Berry phase has accepted that topology can produce phase values within an assumed circle — which removes any worry that the *mechanism* is exotic. The open content is the level-change above it: whether that same topological machinery, pushed down a level, produces the circle itself. That is a real and substantial claim, not a mere matter of how far an established result reaches — but it is a claim about *extending a demonstrated mechanism to a deeper level*, not about inventing an undemonstrated one, and that is the precise and limited sense in which Berry phase makes the proposal less exotic than it first sounds.

---

7. Consequences for the Born-Rule Programme

The memory-residue hypothesis does not derive the Born rule. Probability requires additional structure, and within VERSF that structure is supplied by the ODG and OIP programmes — the companion strands in which the rule attaching probabilities to amplitudes is constructed.

The relation between the two is the actual bridge this paper proposes, and it is worth stating it as a positioning claim rather than a mere division of labour. The Born-rule programme and the memory-residue programme address **opposite sides of the same quantum structure**. The Born-rule programme answers *why probabilities?* — it explains how a probability is extracted from an amplitude, taking the amplitude's phase as given. This paper asks *why phase?* — it proposes an explanation for why amplitudes possess physically meaningful phase structure in the first place, as the continuum representation of accumulated commitment residue rather than as a postulate.

These are complementary questions, and together they reach the two ingredients normally assumed at the foundation of quantum theory:

ODG / OIP → why probability is extracted from amplitudes (probability structure), memory residue → why amplitudes carry phase at all (phase structure).

One explains probability; the other explains phase. The Born-rule programme assumes the phase and derives the probability; the memory-residue hypothesis would explain the phase the Born-rule programme assumes. The two meet on a single shared object — the **amplitude**: the memory-residue programme accounts for its phase, the Born-rule programme for the rule that reads its modulus (squared) as a probability. That shared object is what makes the complementarity structural rather than rhetorical: they are not two unrelated explanations but two accounts of the same complex quantity, one of its argument and one of its magnitude. Together they would furnish *both* foundational ingredients rather than postulating either. This is a contribution to the *interpretation* of the existing programmes, not a replacement for them.

The complementarity can be stated exactly, because the amplitude has exactly two independent real structures and the two programmes address one each.

Proposition 7.1 (The amplitude needs both) (structural; conditional on each programme's inputs)

A quantum amplitude written in polar form

$$\mathcal{A} = A \cdot e^{i\varphi}$$

contains two independent structures: a **magnitude** $A \geq 0$ and a **phase** $\varphi \in U(1)$. These are genuinely independent — neither is recoverable from the other — so a complete account of the amplitude must explain *both*, and an account supplying only one leaves the amplitude half-postulated. The two VERSF strands divide this labour exactly:

- **ODG / OIP** explain the magnitude's physical role: why $|A|^2 = A^2$ is a probability (the Born rule).
- **The Memory Residue Hypothesis** explains the phase's existence: why φ is there at all, as the continuum reading of accumulated commitment residue (§5).

Jointly, the two account for both real degrees of freedom of the complex amplitude — magnitude and phase — rather than postulating either.

What this does and does not claim. It does **not** claim to derive quantum mechanics. The amplitude's polar decomposition into magnitude and phase is itself assumed, as are the state space, the dynamics, and the composition rules; those remain inputs. What the proposition claims is narrower and exact: *of the two real numbers that specify a complex amplitude at the structural level, each is given a physical origin by one of the two programmes, and neither is left as a bare postulate.* That is a statement about how much of the amplitude's *structure* is accounted for, not about deriving the theory built on it. The force of stating it as a proposition is to make visible that the two programmes together cover the complete structural content of the amplitude — A and φ exhaust it — so the joint account is not partial in the dimension it addresses, even though it is partial relative to all of quantum mechanics. A magnitude without a probability rule and a phase without an origin are exactly the two gaps the standard formalism leaves at this level; the two programmes close exactly those two.

---

8. Global Information and Nonlocal Structure

A striking feature of the residue is that it is not localised. No point stores it; no object contains it; it exists only as a property of global topology, measurable only around loops.

This suggests a possible connection to quantum information, which frequently displays global features — entanglement, geometric phase, and topological quantum theories all involve information not reducible to local variables. The memory-residue hypothesis does not explain such phenomena. It identifies a substrate structure possessing the same qualitative character.

One caution keeps this from being a loose "same vibe" gesture, and it is the same distinction §11 draws for Berry phase. The kinship being claimed is *globality* — information carried by loops and relations rather than points — and on that axis the residue genuinely sits with entanglement, geometric phase, and TQFT structure. But the residue is not merely one more global quantity of the familiar kind: what sets it apart, exactly as in §11, is its commitment-origin — it is a property of what was irreversibly committed, not of present global structure that one computes or prepares. So the company it keeps here is real but partial: it shares these structures' globality while differing from all of them in being a record of a fixed past. The claim is that the substrate can support global encoding, not that the residue is the same *kind* of object as a Berry phase or an entangled state.

The significance is therefore architectural rather than explanatory: the substrate already appears capable of supporting globally encoded information — a place for such behaviour to come from, not a derivation of it.

---

9. Why This Matters More for Quantum Theory than Gravity

The implications of the memory-residue hypothesis are asymmetric, and it is worth stating the limit so the proposal does not overreach.

For gravity the impact is modest. The gravity route already proceeds through

commitment → fact density → fact momentum → geometry,

and the memory residue does not modify this chain. It introduces a *parallel* branch from the same root,

commitment → memory residue → phase,

sharing only the origin in commitment. The residue is therefore more closely related to information and phase than to curvature, and its strongest consequences lie on the quantum side of the programme. An indirect link remains conceivable — if geometry emerges from accumulated commitments and those same commitments generate the residue, geometry and memory would be parallel manifestations of one process, one describing quantity and the other history — but this paper claims no more than the parallel.

---

10. The Central Open Question

The entire proposal rests on one unresolved issue, and the preceding sections have read the residue as physics *under the supposition that it survives*. That supposition is exactly what is open.

The closure-side permanence of §3 is established (modulo the inherited essentialness input). What is not established is its transfer to the reconstruction side — to admissible refinement motion, where the residue would have to live to function as a substrate quantity rather than a feature of the closure complex alone. That transfer holds if and only if the primitive-Fact cycle is **refinement-realised**: if some reversible refinement loop transports onto the commitment cycle γ_D.

The decisive question is therefore single and concrete — it is premise (2) of the Interpretation Theorem (§4), isolated from the other two:

Does the primitive-Fact cycle become refinement-realised?

- **If not**, the closure residue is not detected by refinement motion. The memory-residue hypothesis fails as an account of substrate phase, and phase must be interpreted elsewhere.
- **If yes**, the residue survives on the reconstruction side, and — by the persistence result, given the two inputs of §3 — it survives *permanently*. The memory-residue hypothesis becomes a viable account of the origin of phase.

The crux is a category gap: γ_D is a *commitment* object, irreversible and downstream of refinement, while the realising loop would be a *pre-commitment* reversible object. Refinement realisation asks whether admissible transport bridges the two.

It is worth saying where the work to settle this lives, so the question reads as located rather than merely open. It is not settled by anything in this paper, nor by closure topology alone: the closure side already gives γ_D its permanence (§3), so no further closure result decides realisation. What decides it is a property of the admissible-transport map applied to a commitment object — concretely, whether the transport functor of the companion Reversible-Connectedness analysis carries some closed reversible refinement word onto γ_D. That is a question about transport and the merge–split holonomy group, not about phase or commitment as such; it is the same map-theoretic property the parent RC paper isolates as its sole decisive hinge.

The question is decidable in shape, not merely open, and it is worth naming what each verdict would concretely require. A *negative* answer — γ_D unreachable — would hold exactly if the image of the transport map restricted to closed reversible refinement words misses γ_D: if every loop in the merge–split holonomy group transports into the complement of γ_D's class in the closure complex, so that the commitment cycle lies outside the image-from-loops entirely. A *positive* answer requires producing one reversible word whose transport image is γ_D, or proving the image-from-loops contains γ_D's class without exhibiting the word. So the disqualifying condition is definite: $\gamma_D \notin$ image(transport restricted to reversible loops). The reason the question is hard rather than merely unsettled is that this image-from-loops is precisely the object the RC programme has not yet characterised — but it is a characterisation, not a mystery, and the hinge turns on a containment that is in principle checkable once that image is known.

So the question has a definite home (the transport theory of the RC programme), a definite shape (a property of one map on one named cycle), and a definite disqualifying condition (γ_D outside the image-from-loops) — even though it is, as things stand, unsettled.

The question has therefore shifted from *why does phase matter?* to *why does commitment leave a memory?* — and that, in turn, to the one map-theoretic property above.

10.1 Empirical contact

A foundations reader will ask the question this paper has so far not addressed directly: does a memory-residue *origin* of phase differ observationally from a theory in which phase is simply primitive? It is better to answer this plainly than to leave the absence unstated, since an unmentioned gap reads as oversight whereas a stated one reads as discipline.

The honest baseline is that the proposal is *interpretive*, and interpretation does not, by itself, generate new predictions. If the memory-residue origin reproduces exactly the phase structure quantum theory already uses — which is the success condition of §5.3's reproduction claim — then by construction it is observationally equivalent to primitive-phase quantum mechanics in every regime where that reproduction is exact. On that baseline the paper claims an explanatory gain (phase acquires an origin) and no empirical one (no measurement distinguishes the accounts), and that is a coherent and common status for a foundational reinterpretation — the same status the kinetic origin of temperature held before it made contact with fluctuation phenomena.

There are, however, places a genuine signature could appear, and they are worth setting out as *targets* — things that would count as empirical contact if they exist — rather than as claims that they do. All are marked **(open, conjectural)**; none is established here. The value of listing them is to show that the proposal is not interpretive *by necessity*, only interpretive *so far*: it has identifiable places where it could acquire teeth.

The common root of every candidate signature is the same tension the paper has used throughout: the substrate residue is *discrete* ($\mathbb{Z}_7$), while continuum phase is *continuous* (U(1)). §5 reads the discrete residue as the source lighting up a continuous connection — but a discrete source can, in principle, leave a discrete fingerprint somewhere.

This must be reconciled with §5.2, which sold the continuity hard — "the source is discrete, the field it lights up is not." The two are not in tension once the *regime* is fixed. §5.2's continuum holonomy is the description that holds *after* the continuum limit: where emergent time and coarse-graining have smoothed the seven points into a circle, phase is genuinely continuous and no discreteness survives, which is why the proposal is observationally equivalent to primitive-phase QM in all ordinary regimes. The candidate signatures below live in the complementary place: *near the substrate scale, before the continuum limit is taken* — where the smoothing is incomplete and the underlying seven has not yet dissolved. The claim is therefore not that the seven is both visible and invisible, but that it is invisible in the continuum limit (§5.2) and potentially visible only where that limit breaks down (here). A signature, if it exists, is precisely a probe of the regime where the continuum description fails. Four places it might appear:

1. **Regimes where $\mathbb{Z}_7$ discreteness leaks through.** If there is a regime — high closure density, near a commitment event, at the resolution where the continuum description breaks down — in which the underlying seven-valuedness is not yet smoothed into a circle, phase would there take only seventh-root-of-unity values where a primitive continuous phase would allow the full circle. A primitive phase has no reason to quantise in sevens; a commitment-sourced phase would. This is the sharpest possible signature, and the one most directly tied to the paper's central claim.
2. **Constraints on allowed holonomies.** Even where individual values are not resolved, the $\mathbb{Z}_7$ source could constrain which *total* holonomies are admissible — forbidding holonomy values a primitive U(1) theory would permit, or favouring multiples of $2\pi/7$. A measured spectrum of allowed phases that clustered on or avoided the seventh roots would be evidence; a featureless continuum would not falsify the proposal (it is the expected coarse-grained limit) but would leave it interpretive.

3. **Finite-resolution phase structure.** If phase is the continuum shadow of a finite discrete structure, there may be a fundamental limit to phase resolution — a smallest meaningful phase difference set by the seven-step substrate, below which the continuum description has no referent. Primitive-phase theory posits an exactly continuous phase with no such floor. Detecting a resolution floor of the right structure would distinguish them.
4. **Topological memory effects.** If commitment genuinely leaves a persistent residue, there may be observable consequences of the residue's *path-dependence and non-locality* — interference outcomes that depend on enclosed committed structure in a way a purely local-phase account cannot reproduce, analogous to but distinct from the Aharonov–Bohm effect, with committed regions playing the role of enclosed flux.

None of these is claimed to occur, and the honest baseline above stands: if the reproduction of §5.3 is exact in all accessible regimes, the proposal is observationally equivalent to primitive-phase quantum mechanics and the targets are empty. But the targets are not empty *by construction* — each names a specific way the discrete substrate could fail to be perfectly hidden, and any one of them, if realised, would convert the proposal from an interpretation into a testable hypothesis. That is the highest-value development available to the programme, and it is the reason the $\mathbb{Z}_7$ structure is worth taking seriously as physics rather than as bookkeeping: a primitive phase is continuous because it is assumed to be; a commitment-sourced phase is continuous only in the limit, and limits can be probed.

---

11. Why Physics Does Not Already Have a History Variable

It would be easy to mistake the memory residue for one more entry in the long list of conserved quantities. It is worth saying directly why it is not — why it belongs to a category physics does not currently contain — because that, more than any single consequence, is what would make the proposal matter.

Consider the quantities physics already has. Every one of them is a function of the present state.

- **Energy** is state-dependent: it is read from the current configuration and its rate of change. Two systems in identical present states have identical energy, whatever their pasts.
- **Momentum** is state-dependent: it is fixed by what is moving how, now.
- **Charge** is state-dependent: it is a property of what is currently present.
- **Entropy** is state-dependent: it is a count over the current macrostate. It is often *called* a quantity that "knows about time," but it does not. Entropy is a function of the present coarse-grained state; the second law describes how that function tends to change, but the value itself carries no record of which past produced it. Two systems in the same macrostate have the same entropy regardless of how they arrived.

The pattern is exact and worth naming. Every standard physical quantity is recoverable, in principle, from a sufficiently complete description of the present instant. Time enters physics only through *equations of motion* — rules relating the present to the next instant — never through a *variable* whose value is itself a fact about the past. The state is sufficient; history is reconstructed, when it is reconstructed at all, by integrating dynamics backward, not by reading a quantity that stored it.

The memory residue would break this pattern. It is not state-dependent; it is **commitment-dependent**. Its value is not a function of the present configuration — two substrates in identical present states could carry different residues if different nets of commitment had occurred to produce them. The residue is not recoverable from the instantaneous state, because it is not *in* the instantaneous state; it is in the global topology that irreversible commitment has shaped, and that topology is a fact about what has been committed, not about what currently is.

Two near-cousins will occur to any physicist at this point, and naming the difference in each case hardens the claim rather than weakening it.

Hysteresis and path-dependent state variables — magnetisation, plastic strain, spin-glass configuration — look like memory: their present value depends on the path taken to reach it. But that "memory" is fully reducible to present microconfiguration. The magnetic domains are physically *there* now; read the current microstate completely and the history-dependence vanishes, because what looked like a record of the past is just a present arrangement that a coarse description failed to capture. These remain state functions at the micro level. The memory residue is categorically different precisely because it is *not* reducible to any present configuration, fine-grained or not: there is no microstate reading that recovers it, because it is not stored in the state at all.

Geometric (Berry) phase is the closest existing thing to a history variable, and the paper leans on it twice: in §6.1 as the proof-of-concept that phase *can* be topological rather than primitive, and here as the nearest cousin to a history variable. It is genuinely route-dependent and genuinely non-local. But it too is computed from present-specified data — the present connection, the present Hamiltonian path — not from what was irreversibly committed. A Berry phase is a property of a transport one *chooses* to perform on a present system; the memory residue is a property of transport reality has already *committed*, and cannot un-commit. So Berry phase plays a double role for the proposal: it shows the mechanism (topology yielding phase) is real and established, while marking the exact line the hypothesis crosses beyond it (from chosen present transport to fixed past commitment). Same holonomic form, opposite ontological status — one read off present structure, the other off a fixed past.

This is the qualitative difference, and it is a difference of category, not of magnitude:

energy, momentum, charge, entropy → functions of the present state, hysteretic / Berry-type quantities → functions of present (micro)state or chosen transport, memory residue → a function of accumulated commitment, irreducible to any present state.

A new conserved quantity would be a new *state* function — another thing readable from the present. The memory residue would be something physics has not had: a quantity whose very existence presupposes that the past is fixed, and whose value records that fixing. It is conserved not in the way charge is conserved (because the present cannot create or destroy it) but in the way the past is conserved (because what has been committed cannot be uncommitted).

The distinction is not semantic; it changes the ontology. Every quantity now in physics is a property of *what exists* — of the configuration present at an instant. The memory residue would be a property of *what happened* — of the commitments that have occurred, whether or not any trace of them remains in the present configuration. The claim, stated at its defensible strength, is this: *no quantity in established fundamental physics is a history variable in this sense.* That is why the residue is not merely another invariant to add to the list. It would be the first member of a category the list does not contain, and the rest of this paper's interest follows from that categorical novelty, not from the residue's size, sign, or specific value.

---

12. The First History Variable

Should the open question of §10 resolve affirmatively, the residue would occupy exactly the empty category §11 identifies.

Most physical quantities describe the present. Position describes where something is; momentum, how it moves; energy, what transformations remain possible. The memory residue would be different: its value depends upon prior commitment. It is therefore neither purely geometric nor purely dynamical, but *historical*.

What kind of history must be stated precisely, because it is far weaker than the word usually implies. Not chronology. Not narrative. Not a record of any particular event. Only the net, cancelling tally — valued in $\mathbb{Z}_7$, certifying *that* commitment occurred, blind to what, where, when, and even to how many beyond the count mod seven.

Even at that strength, this would be the first genuine **candidate fundamental history variable** identified within VERSF — the programme's first candidate quantity that would exist because the past is fixed rather than because of any configuration in the present, and (if the framework holds) the first in physics. The qualifier "candidate" is deliberate and is dropped only when the open question of §10 is closed: until refinement realisation is settled, the residue is a history variable *in waiting* — established on the closure side, conditional on the reconstruction side — and calling it the first *candidate* loses nothing while remaining exact. The universe would retain information unavailable from the instantaneous state alone: the present would depend not only on current structure but on accumulated commitment. What it stores is not nothing — it is exactly a bounded topological tally, a few bits per protected cycle (§4) — but it stores *no chronological or narrative information*, only that bounded count. The contrast with ordinary records is therefore not "information versus none" but a bounded, history-only tally versus an open-ended archive; and that bounded budget is itself the feature that makes it a witness rather than a diary. It would constitute a genuinely new category of physical quantity.

It is a closure-side reality already (given essentialness); it becomes a variable of the reconstruction substrate the moment refinement realisation is established, and not before.

---

13. Programme Consequences

Before drawing the consequences, it is worth observing that phase was, in retrospect, the natural place this programme was always heading — that the route to it is not an accident but a chain in which each step forces the next. Read forward, the chain is

commitment → protected cycles → holonomy → phase,

and each arrow is established earlier in this paper:

1. **Commitment creates irreversibility.** A Fact is, by definition, an event that cannot be undone (§1, §3).
2. **Irreversibility plus essentialness creates protected cycles.** A loop around an essential discarded region cannot be contracted — not through the region (irreversibility) nor around it (essentialness) — so commitment manufactures permanently non-bounding cycles, *given essentialness* (§3). The essentialness input is inherited, not commitment-native; the chain produces protected cycles only when commitment isolates an essential region, exactly as §3 was careful to isolate.
3. **Protected cycles create holonomy.** A non-bounding cycle is a nontrivial first-homology class; a charge carried around it is a holonomy, global and route-dependent by construction (§6, §6.0b).
4. **Holonomy has the phase signature.** That holonomy possesses, one for one, the four defining properties of phase (§6.1, Prop 6.1), and — given emergent time — covers phase as such (§5.3).

Read as a whole, the chain says that *given commitment and the essentialness of what it discards, something with the structure of phase is the expected terminus*, not a surprise stumbled upon. Once a framework takes irreversible commitment as primitive and its discarded regions as essential, it is led — through protected cycles and their holonomy — to a global, route-dependent, loop-accumulated, interference-bearing quantity, which is precisely what phase is. The Memory Residue Hypothesis therefore reads less as a bold conjecture bolted onto the programme than as the step the programme was always going to reach: the place the logic of commitment points once one asks what irreversibility leaves behind. (The chain establishes the *destination*, not that the journey completes — it rests on the inherited essentialness input of step 1, and the single open hinge of §10 still gates whether the residue is realised; what is inevitable is where the residue must live and what it must look like, not that it is nonzero.)

With that framing, the consequences. If the memory-residue hypothesis survives further scrutiny, they extend well beyond the Gate-3 result that prompted it, and they are the reason the open question of §10 is worth the effort of settling.

First, the interpretation of phase changes. Phase ceases to be primitive and becomes inherited. The most basic structure of quantum amplitudes — that they carry a phase at all — would no longer be a starting assumption but a consequence of commitment leaving a topological trace. This is phase *as such*, not merely its geometric sector: the extension turns on emergent time leaving no separate fundamental dynamical phase (§5.3). One step in that extension is itself open — whether flat-sector holonomy reproduces the effective dynamical phase, not merely displaces a primitive — and is marked so; granting it, what the substrate offers is the whole of phase, with the geometric/dynamical split recovered only downstream.

Second, the role of commitment changes. Within the existing programme, commitment generates facts and, through fact density and fact momentum, geometry. Under the hypothesis it does something further: it generates *memory*. Commitment becomes the source not only of what is, but of a persistent record that something was.

Third, the ontology of the framework changes. The substrate would contain not only structure but history — a quantity (§11–12) that is a property of what happened rather than of what exists. The framework would then describe a reality that carries its own past in its shape, not merely in the records that past leaves behind.

Fourth, the quantum and closure programmes become more closely linked. The same commitment events responsible for closure residue would be responsible for the phase structure quantum theory requires. Two strands developed largely independently — closure topology on one side, amplitude and phase on the other — would turn out to draw on a single underlying source.

Finally, the programme acquires a new explanatory direction. Alongside the established gravity route,

commitment → fact density → fact momentum → geometry,

the hypothesis opens a second, parallel route,

commitment → memory → phase → quantum structure,

independent of the first and sharing only its origin in commitment. The memory-residue hypothesis therefore does not merely add an interpretation to an existing result; it opens a second major branch of the framework, running from the same root toward quantum structure rather than toward geometry. That is the sense in which the open question of §10 is not a local technicality but a fork in the programme's development.

---

14. Conclusion

The memory-residue hypothesis proposes that persistent closure holonomy supplies the physical origin of phase.

The proposal is conditional. It depends on the persistence branch of the Gate-3 programme — which itself rests on two inputs, irreversibility (native) and essentialness of the discarded region (inherited) — and, for transfer to the substrate, on the open question of refinement realisation. None of this is claimed as settled, and the Interpretation Theorem of §4 states exactly which three inputs the conclusion requires.

Nevertheless the implications are substantial. A surviving residue would transform phase from a primitive structure into an inherited one, connecting commitment, topology, transport, and phase within a single framework through the chain

$K = 7$ closure geometry → $\mathbb{Z}_7$ transport residue → closure holonomy → U(1) phase,

discrete part of the final embedding elementary, continuum part inherited from Fact Momentum. And it would introduce a kind of physical quantity the existing inventory does not contain — not another function of the present state, but the first quantity whose existence presupposes a fixed past:

Not a quantity describing what exists. A quantity describing what happened. A history variable.

If correct, the usual order of explanation inverts. We are accustomed to thinking that the past is available to us because something survived to record it — that memory is downstream of records, and records downstream of events. The memory-residue picture turns that around. The records are not the source of the past's persistence; they are a consequence of it. What makes the past fixed is not that something happened to be written down, but that commitment, being irreversible, leaves a permanent mark in the shape of reality itself — a mark no later motion can reach back and erase.

There is a precedent for what this kind of move accomplishes. Temperature was, for a long time, a primitive — a reading on an instrument, governed by laws that described how it behaved but never said what it *was*. Thermodynamics eventually supplied a physical origin: temperature is mean molecular motion, and the gas laws became consequences of mechanics rather than axioms about an unexplained quantity. If the memory-residue hypothesis survives, quantum theory would acquire something of the same kind — a physical origin for phase, where there is now only an inherited mathematical ingredient. Phase would cease to be the thermometer reading no one can account for and become the holonomy of accumulated commitment, with the formalism's phase factors as consequences rather than postulates. The analogy is not that the mathematics resembles thermodynamics; it is that the *explanatory step* is the same — turning a primitive into a derived quantity by identifying what it is made of.

And the identification would reach all the way down to the programme's deepest structural claim. The memory residue counts in sevens because the closure architecture counts in sevens. The appearance of phase would then be the continuum manifestation of the same $K = 7$ closure structure that governs the substrate's admissible transport sector — periodicity of phase inheriting, through emergent time and the continuum limit, the sevenfold periodicity of the substrate itself. What looks in the continuum like a featureless circle would be, at bottom, the shadow of a discrete seven.

The universe does not remember because records survive.

Records survive because commitment leaves topology behind.