The State of the VERSF Programme

From Distinction, Closure, and Admissibility to Geometry, Matter, and the Final Substrate Questions

Keith Taylor

VERSF Theoretical Physics Programme

A note on how to read this document. This is a review of where the VERSF programme currently stands. It is written on two levels at once. Each major idea is first stated in the programme's own technical language, and then re-stated *in plain terms* for readers approaching the work for the first time. The plain-language passages are not simplifications that change the claims — they are translations of them. Where a result is secure it is marked **proven**; where it depends on assumptions still being tested it is marked **conditional**; where it is a working hypothesis it is marked **conjectural**. Nothing in the technical spine has been softened to make the plain-language layer read more smoothly.

Abstract

The VERSF programme began with a small collection of foundational principles: finite distinguishability, irreversible commitment, closure, admissibility, and emergent time. Over a sequence of interconnected papers these principles have been developed into a coherent emergence framework that reconstructs large portions of modern physical structure from substrate-level assumptions.

The programme has progressively reduced apparent multiplicity to structural necessity. Questions that initially appeared independent — probability, Hilbert geometry, gauge structure, Lorentzian geometry, transport, matter classification, and gravitational dynamics — have increasingly been shown to arise from common underlying admissibility constraints.

This paper reviews the current status of the programme. We summarize which components are substantially closed, which remain conditional, and which open questions now define the frontier of the framework. We argue that the programme has reached a stage where most major structural layers have been constructed, and that the remaining unresolved issues have compressed into a small number of sharply defined substrate questions.

In plain terms. VERSF starts from an unusually small set of starting assumptions and tries to *grow* the rest of physics out of them, rather than writing the physics in by hand. The surprise of

the programme so far is not that the assumptions are powerful — many frameworks make powerful assumptions — but that things which look like separate problems (why probabilities work the way they do, why space has the shape it has, why there are forces and particles) keep turning out to be the same problem wearing different clothes. This document is a stock-take: what looks solid, what is still resting on unverified scaffolding, and what is genuinely unknown. The headline is that the list of genuinely unknown things has become very short.

Contents

---

1. Introduction

The central ambition of VERSF has always been reconstruction rather than postulation.

Rather than assuming geometry, fields, particles, probability, or an emergent geometric continuum, the programme asks:

What is the minimum structure required for distinguishable physical reality to exist?

The resulting framework begins not with a continuum, matter, or quantum theory, but with:

- finite distinguishability,
- admissible comparison,
- closure,
- commitment,
- record formation,
- refinement.

The programme has pursued two parallel goals:

1. Reconstruction of known physics.
2. Identification of unavoidable residual structures implied by the substrate.

The first goal has progressed significantly.

The second has gradually reduced to a remarkably small set of remaining questions.

In plain terms. Most physics takes the stage and props as given — there is space, there is time, there are particles — and then writes down rules for how the props move. VERSF refuses to put the props on the stage in advance. Instead it asks the most stripped-down question it can: *if reality is going to contain things that can be told apart from one another at all, what is the least it must contain?* The answer it works from is a handful of primitives — the ability to tell two things apart (finite distinguishability), the ability to compare them in a legitimate way (admissible comparison), the way comparisons can loop back on themselves (closure), the moment a possibility becomes a settled fact (commitment), the leaving of a trace (record formation), and the ability to look more finely (refinement). From this short list the programme has two jobs: rebuild the physics we already know, and find out what *else* the list secretly forces to exist whether we wanted it or not. The first job is mostly done. The second job has narrowed down to almost nothing — and that "almost nothing" is what the rest of this paper is about.

---

2. What Has Been Substantially Established

2.1 Emergent Time

One of the earliest conclusions of the programme is that time is not fundamental.

The framework distinguishes:

- substrate ticks,
- commitment events,
- committed records,
- physical time.

Physical time emerges from the accumulation of committed facts rather than existing as a primitive coordinate. **(proven within the framework.)**

This shift underlies every subsequent construction.

In plain terms. In everyday physics, time is a coordinate — a number line you slide along. VERSF says that picture is downstream of something more basic. What actually happens at the bottom level is that possibilities get *settled* — a commitment event occurs, a fact is recorded, and that record cannot be un-made. Pile up enough of those irreversible settlings and you get something that behaves like the flow of time. Time, on this view, is not the river you float down; it is the silt left behind as facts accumulate. The reason this matters is that everything built later in the programme inherits it: there is no background clock to lean on, so every construction has to earn its sense of "before and after" from records, not from a pre-drawn timeline.

2.2 Operational Geometry

A major milestone was the emergence of operational geometry.

A sequence of papers established:

- operational distance,
- operational curvature,
- operational projection,
- operational inner products,
- refinement geometry.

These culminated in Unified Refinement Hilbert Geometry, where geometry, projection, and probability emerge from a common admissibility structure. **(proven within the framework, conditional on the admissibility axioms.)**

The programme therefore no longer treats geometry and quantum structure as independent domains.

In plain terms. "Geometry" here does not start with rulers and angles. It starts with a question: *how hard is it to tell two states apart?* If two configurations are nearly indistinguishable, call them close; if they are easy to separate, call them far. Define distance that way and a genuine notion of shape appears — including curvature, the same quantity that in ordinary physics tells you a surface is bent. The striking part is the destination. When the programme followed this "distinguishability density" idea all the way through, the structure it landed on was the same mathematical machinery quantum mechanics uses — inner products, projections, the geometry of a Hilbert space. So the framework stopped treating "the shape of space" and "the rules of quantum states" as two subjects. They are one subject seen from two angles, both falling out of which comparisons are *allowed* (admissible).

2.3 Probability and the Born Rule

Multiple independent derivations converged on the same result:

Probability must scale quadratically. **(proven within the framework, conditional on reversible-mixing and admissible-transport assumptions.)**

The development proceeded through:

- admissibility measures,
- operational distinguishability geometry,
- universal measure arguments,
- squaring-residue analysis.

The programme ultimately reduced the Born-rule question to a small number of structural assumptions concerning reversible mixing and admissible transport.

This represents one of the strongest convergence points in the entire framework.

In plain terms. One of the oldest puzzles in quantum theory is the Born rule: to get the probability of an outcome you take the amplitude and *square* it. Why squared? Why not the amplitude itself, or its cube? In standard quantum mechanics this is essentially put in by hand because it matches experiment. VERSF tried to *derive* the squaring — and not once, but along several independent routes that did not know about each other. They all arrived at the same verdict: given the way distinguishability and legitimate mixing work in the framework, probability has no choice but to scale quadratically. When several unrelated arguments march to the same answer, that answer is hard to dismiss as an accident of one particular method. That convergence — written schematically as the rule that weights go as the square, $P \propto |\cdot|^2$ — is one of the strongest pillars the programme has.

2.4 Gauge Structure

Persistent transport was shown to survive refinement through cohomological invariants.

This led naturally to:

- gauge potentials,
- Wilson loops,
- gauge transport,
- source currents,
- matter coupling.

The resulting framework reconstructs a Maxwell-like gauge sector without assuming gauge symmetry as fundamental. **(conditional.)**

In plain terms. Forces like electromagnetism are described by "gauge structure" — a precise bookkeeping of how a quantity gets carried around a loop and whether it comes back changed. VERSF did not assume any of this. It started from the humbler observation that *some* relationships survive when you look more finely (they are robust under refinement), and asked what survives and what washes out. The things that survive turn out to be exactly the quantities a physicist recognizes as gauge potentials and Wilson loops — the carry-it-around-a-loop bookkeeping. From there a Maxwell-like force sector appears, complete with sources and a way for matter to feel it. The point worth stressing: the gauge symmetry usually treated as a fundamental input is here an *output*, a leftover of which transports refuse to be erased.

2.5 Lorentz Geometry and Gravity

A substantial body of work established:

- causal transport,
- transport cones,
- Lorentzian structure,
- metric emergence,
- stress-energy emergence,
- Einstein-type dynamics,
- Einstein-Hilbert emergence.

The gravity programme is now largely structurally complete. **(conditional; internal derivation chain largely present.)**

Open questions remain concerning empirical validation and uniqueness, but the internal derivation chain is largely present.

In plain terms. Gravity, in Einstein's account, is geometry: matter and energy tell the geometric structure how to curve, and the curvature tells matter how to move. VERSF rebuilds this without starting from a metric. It begins with which influences can reach which — a notion of cause and effect (causal transport) that carves out "cones" of what is reachable. Those cones are the seed of Lorentzian structure: the asymmetry between past-and-future directions and sideways directions that gives the continuum its causal character. Out of that grows a metric, a stress-energy content, and finally dynamics of the Einstein type — the same form as the equations of general relativity, including the action principle (Einstein–Hilbert) they descend from. The internal logic is now largely in place. What is *not* settled is whether this reconstruction is the only one possible (uniqueness) and how it lines up against measurement (empirical validation). So the chain is built; the audit against the real world is the remaining work.

2.6 Matter and Particle Structure

The programme has also developed:

- Persistent Fold Defects,
- closure-norm condensation,
- substrate stiffness hierarchies,
- finite distinguishability spectra,
- charged-lepton reconstruction.

These provide a candidate route from substrate defects to Standard Model particle classes. **(conjectural to conditional; least mature sector.)**

The matter sector remains less mature than the gravity sector, but its overall architecture is now visible.

In plain terms. If geometry and forces come from the substrate, particles ought to as well — and the candidate answer is that particles are *defects*: stable knots or snags in the fold structure that refuse to relax away (Persistent Fold Defects). Different defects, stabilized by different "stiffnesses" of the substrate, would correspond to different kinds of particle, and because distinguishability is finite, only a finite menu of stable configurations is allowed (finite distinguishability spectra). The programme has gone furthest with the charged leptons — the electron and its heavier cousins — sketching how their pattern could arise this way. Honesty requires flagging that this is the youngest and least secure part of the framework. The shape of the explanation is visible; the explanation itself is not yet locked down. Treat this section as a promising blueprint, not a finished building.

---

3. The K = 7 Closure Architecture

A recurring feature of the programme is the appearance of the K = 7 closure structure.

The framework now contains:

- explicit K = 7 closure graphs,
- spectral analysis,
- transport groups,
- closure-frame constructions,
- Wilson-loop structures,
- closure holonomy sectors.

The transport-group analysis identifies:

$$G \cong D_7$$

with

$$\ker(\pi\_or) \cong \mathbb{Z}_7.$$

This establishes a genuine sevenfold closure channel within transport. **(proven within the framework.)**

The remaining question is not whether the channel exists algebraically.

The question is whether the substrate populates it physically — and that remains open. It turns on a single property of how committed outcomes are read out of the substrate, taken up in Section 7. **(open.)**

In plain terms. A particular number keeps showing up unbidden: seven. When the programme analyzes how transports loop and close on themselves, the symmetry group that governs them is D_7 — the symmetry of a seven-sided figure, the dihedral group on 7 elements. Hidden inside it (technically, the part that the orientation map sends to nothing — its kernel) is $\mathbb{Z}_7$, the clean sevenfold cyclic symmetry. In plain language: there is a real, mathematically forced "sevenfold channel" running through the structure of how things can be carried around and brought back. That much is settled. But a channel existing on paper is different from the channel being *used*. The open question — the one the rest of the paper builds toward — is whether nature actually pours anything down this sevenfold channel, or whether it sits there empty. It is worth saying clearly what the seven *is* here: not a clock or a phase that ticks round, but a seven-mark *alphabet* for comparing neighbouring sites — and, deeper down, the number of times a closure loop wraps before it closes. Whether the channel is occupied hinges on a single question about how the substrate's records are read, which Section 7 makes precise.

---

4. Gate 3: The Final Structural Residue

The Gate-3 programme has undergone substantial compression.

Originally framed as a broad question concerning residual transport structure, it has progressively reduced through:

- locality analysis,
- reachability analysis,
- transport-group analysis,
- closure-connection construction,
- topology analysis,
- operational-quotient analysis.

The result is a sharp formulation.

The existence question becomes:

$A \subseteq B^1$?

where:

- A is the admissible offset space,
- B^1 is the coboundary space.

If $A \subseteq B^1$, every admissible configuration is pure relabelling and the closure sector collapses.

If $A \not\subseteq B^1$, nontrivial closure holonomy survives.

This is the primary existence question, and it remains open. What has been established is its precise form rather than its answer. The closure charge is a *contextuality obstruction* — a class in the first cohomology of the vacuum transport complex, nonzero exactly when the local pairwise comparisons cannot be reconciled into a single consistent global labelling of the sites. (Two further things about this class are settled separately: in which group the *measurable* charge is valued — its register — is fixed in Section 6, and turns out to be the two-valued class, not the seven-valued one; the seven persists as the comparison alphabet and as the transport-group kernel, not as the register.) Whether the comparisons reconcile — whether $A \subseteq B^1$ — has been reduced to one property of the readout rule (Section 7), and is not settled by local definiteness alone: a configuration can be perfectly definite on every edge yet globally frustrated, like an impossible staircase whose every step is fine but whose full loop returns higher than it began. **(open.)**

In plain terms. "Gate 3" started life as a vague worry — *is there some leftover structure in the substrate we haven't accounted for?* — and the programme has spent a long time squeezing that worry down until it became a single yes-or-no question. Here is the question without the machinery. Picture a web of sites, with each connecting edge carrying a comparison label that says how its two ends differ. The question is whether all these local comparisons can be reconciled into one consistent global numbering of every site — like checking whether "A is 3 ahead of B, B is 2 ahead of C" statements can be satisfied by giving each site a single position — or whether they are *globally frustrated*, like an impossible staircase where every individual step looks fine but going all the way around brings you back higher than you started.

Is A entirely contained in B^1?

If yes ($A \subseteq B^1$), there is no hidden structure — every apparent difference is just relabelling, and the closure sector collapses into nothing. If no ($A \not\subseteq B^1$), something real survives that cannot be relabelled away — a genuine, irreducible loop-memory the programme calls closure holonomy. The crucial point, now proven, is that local definiteness does not decide this: you can have a web where every single comparison is perfectly sharp, yet no global numbering reconciles them. So the question is not whether the comparisons are definite — they are — but whether the framework's rules *force* a global reconciliation or merely permit local consistency. (Section 6 settles a different question about the same charge: not whether it is there, but what kind of thing it is if it is — and the answer is a two-way switch, not a seven-position dial.)

---

5. Fact Momentum and the Continuum Limit

A major recent result is the identification of the relationship between Gate 3 and Fact Momentum.

The programme now supports the theorem:

Gate-3 closure holonomy is the flat sector of the continuum connection whose stress-energy is Fact Momentum. **(proven within the framework, conditional on the continuum-limit construction.)**

Fact Momentum decomposes into:

- a curvature sector,
- a flat holonomy sector.

The curvature sector carries local field energy.

The flat sector carries topological memory.

Gate 3 is exactly this flat sector.

In a flat limit:

Fact Momentum = Gate 3.

In the generic case:

Gate 3 ⊆ Fact Momentum (the topological summand).

This places the closure programme naturally inside the continuum theory.

In plain terms. This is the result that ties the two halves of the programme together. "Fact Momentum" is the framework's name for the energy-and-momentum content that the continuum carries — its version of the stress-energy that, in gravity, tells geometry how to bend. The new theorem says Fact Momentum splits cleanly into two parts. One part is *curvature*: ordinary, local field energy, the kind that bends things in your immediate neighbourhood. The other part is *flat*: it carries no local bending at all, yet it is not nothing — it stores a kind of *topological memory*, information about how loops wrap around globally that no purely local measurement can detect. The claim is that Gate 3 — the sevenfold closure leftover from Section 4 — is precisely this flat, memory-carrying part. In the special case where there is no local curvature at all, the entire Fact Momentum *is* Gate 3. In general, Gate 3 is the topological slice of it. The payoff is conceptual: the closure question is no longer a side-quest. It lives inside the main continuum theory as one of its two natural pieces.

---

6. The Register of the Charge

A question the arc carried as a fixed assumption rather than a result can now be answered: granting that a closure charge survives, *what kind of charge is it* — what group does it take values in? The earlier arc wrote the class with seven-valued coefficients, the seven inherited from the K = 7 constraint count used as a comparison alphabet. The foundational theorems of the programme settle the matter, and the answer is that the *measurable* charge is the two-valued ($\mathbb{Z}_2$) class, not the seven-valued one.

This is independent of occupancy. Whether or not the sector turns out to be occupied (Section 7), the register question asks only: *if* a charge is physically there, what is it valued in? The answer holds on either occupancy branch.

The reduction. The whole holonomy content of an edge label lives in one place. Any admissible offset splits as $\rho_{uv} = \chi(v) - \chi(u) + \tau_{uv}$, where the χ-part is a pure coboundary — a relabelling — and τ is the fold/interface transition. The χ-part telescopes: its sum around any closed loop vanishes identically. So on every cycle γ the loop sum Σ_γ ρ ≡ Σ_γ τ. The charge is therefore the register of τ, the measurable transition an interface records on transport around a loop, and the register question is exactly: *what can a measurable τ be valued in?*

The first floor — uniqueness. The minimal unit of distinguishability is the fold: a two-state object whose intrinsic symmetry is $\mathbb{Z}_2$, and this is forced, not chosen — the reversible transformations on one bit form the unique two-element group, a fact the framework's own robustness audit marks as *impossible* to be otherwise. Moreover every operationally accessible distinction factors through binary folds: there is no admissible measurement whose answer space is irreducibly seven-valued and does not decompose into binary partitions. A measurable seven-valued register is therefore not an admissible operational primitive. **(proven, given the Fold Uniqueness Theorem and the operational reading of the charge.)**

The second floor — saturation. This closes the one gap uniqueness leaves: could a seven-valued register arise not as a primitive but as a composite or emergent observable? No. Once a theory contains the fold, the entire algebra of observables is the one the fold generates — the 4×4 complex matrix algebra $M_4(\mathbb{C})$ over the fold's state space — with nothing left over. That algebra carries no seven-valued grading; the finite structure it does carry is the $\mathbb{Z}_2$ of orientation and the $\mathbb{Z}_3$ of the colour split, never a $\mathbb{Z}_7$. A measurable seven-valued holonomy would be an observable outside the fold algebra — precisely the falsifier the saturation theorem names for its own refutation. **(proven, given the Saturation Theorem and the operational reading.)**

The verdict. Assembling the reduction and the two floors: the observable Gate-3 closure charge cannot be seven-valued. The register available — the $\mathbb{Z}_2$ of the fold automorphism, of orientation, of fact-trapping, of the spin closure phase — is what it must be:

$$\kappa \in H^1(\Gamma\_vac; \mathbb{Z}_2).$$

When occupied, it is occupied by the same trapping argument that makes facts necessary in the first place, and it unifies with the binary registers the rest of the corpus realizes throughout. **(proven, conditional only on the operational reading; independent of the occupancy branch.)**

Where the seven actually lives. None of this erases the seven. It keeps two legitimate roles and loses a third. It remains the kernel of the transport group — $\ker(\pi\_or) \cong \mathbb{Z}_7$ inside D_7, the genuine sevenfold channel of Section 3 — and it remains a *count*: of admissibility constraints ($K = 7$), of loop channels ($N\_loop = 14 = 2K$), of gauge generators ($8 + 3 + 1 = 12$). What it is not, for any *observable* charge, is the order of a register an experiment could read. The earlier seven-valued cohomology class was never wrong as cochain bookkeeping on the fixed complex; it simply cannot be the register of the measurable holonomy, because the observable sector admits no seven-valued register at all.

Security. This verdict is, unusually, *more* secure than the occupancy question it sits beside. Occupancy waits on an axiom whose reading is not yet fixed; the register verdict rests on the framework's most secure tier — the uniqueness of the minimal distinction and the four-dimensionality of the fold state space, both Category I results the audit marks unconditional. A reader who accepts the gauge group, the fine-structure constant, or the cosmological constant has already accepted the harder results the register verdict needs.

In plain terms. Suppose the sevenfold channel *is* occupied and carries a real charge. What kind of charge is it — a seven-position dial, or a simple two-way switch? The answer, forced by the framework's own foundations, is the switch. Here is why, in two moves. First, the entire charge collapses onto one ingredient: the comparison a fold-interface records as you carry it around a loop. Second, that interface is a *fold* — the framework's minimal unit of distinction — and a fold is irreducibly two-valued. Two theorems pin this down. One says the smallest possible distinction is a single yes/no, and that *every* measurement you could ever make breaks down into a sequence of yes/no folds — so there is no genuine seven-way measurement hiding anywhere. The other says that once you have the fold, the whole catalogue of measurable quantities is fixed and contains no seven-valued slot for such a charge to occupy. A measurable seven-valued charge would not just be odd — it would break the theorem that says the fold is enough. The seven hasn't vanished; it is doing a different job than a dial. It *counts* things — seven independent constraints, fourteen loop channels, twelve force-carriers — and it labels the sevenfold transport channel itself. What it never does is become a seven-position reading on a meter. The measurable charge is the two-way switch: $\mathbb{Z}_2$. And this is one of the safest conclusions in the whole framework — safer than the charge being there at all, because it leans only on the most bedrock results.

---

7. What Remains Open

The open questions are these. Occupancy rests on one property of the readout rule; observability sits behind it, arising only if occupancy holds; and a single seam in the register verdict of Section 6 remains genuinely open, concerning a layer no measurement reaches.

7.1 Occupancy — The Readout Fork

Does admissibility populate the closure sector? Equivalently:

$A \subseteq B^1$?

This is open. What the arc has done is reduce it — exactly, not heuristically — to a single property of one principle.

Several structural conditions are discharged. Energy conservation forces every admissible offset to be *closed* ($A \subseteq Z^1$) — a derived flatness, not an assumption. The topology determines whether a nontrivial first cohomology is even available (the T1 branch); on the trivial branch $\kappa = 0$ automatically and the question is vacuous. And a survival criterion fuses the two: given availability, the sector is occupied precisely when admissibility permits a closed offset that is not exact. None of these forces $A \subseteq B^1$.

What is *not* settled is whether the remaining principle — Uniform Readout, the rule governing how committed outcomes are extracted — forces a global reconciliation. An equivalence theorem makes this exact: Uniform Readout forces the charge to vanish if and only if reading out the substrate amounts to consulting a single, path-independent global record — a "ledger" whose value at each site does not depend on the route taken to reach it. This is the same distinction the measurement literature draws between merely-local data and globally-reconcilable data. The question therefore reduces to a fork in what Uniform Readout asserts:

- **UR-frame** — readout refers every site to a single global frame; commitment lays down (or readout presupposes) a global, path-independent ledger. Then $A \subseteq B^1$, $\kappa = 0$, and the sector is empty (Branch A).
- **UR-comp** — readout applies a uniform comparison rule on each edge but refers outcomes only to local frames; commitment lays down only local comparisons. Then $A \not\subset B^1$ is permitted, and the sector can be occupied on a T1 topology (Branch B).

(open — the entire occupancy question now rests on which of these two the Uniform Readout axiom asserts.)

The arc does not claim to know which. It records one asymmetry, marked as parsimony rather than proof: every primitive in the generative chain — Ticks → Fold → Bit → Fact → Record — is *local*, so no primitive operation directly lays down a global frame. A path-independent ledger could only arise as an aggregate of local commitments that happen to be jointly reconcilable, which a T1 topology does not force. Branch A therefore requires the *presence* of a global-frame-fixing mechanism the local ontology does not supply; Branch B requires its *absence* plus an available topology. That locates the burden on Branch A; it does not decide the fork. **(parsimony asymmetry, conditional on the local-ontology premise — not a decision.)**

In plain terms. This is the bottleneck question, and it is open. The closure charge is not a leftover from any smooth "phase" and not a clock that ticks round. It is a property of how a web of committed comparisons fits together. Each edge carries a comparison label; the charge is nonzero exactly when those local comparisons *cannot* be reconciled into one consistent global numbering of the sites — the impossible-staircase situation. The key proven fact is that local definiteness does not settle this either way: a web can be perfectly definite edge-by-edge and still be globally frustrated. So everything comes down to one property of how the substrate is read out. If reading the substrate means consulting a single global ledger whose value at each site is fixed regardless of the path you took to reach it, then the comparisons are automatically reconcilable, the charge vanishes, and the channel is empty. If instead the substrate only ever commits *local* comparisons between neighbours — which fits the framework's picture of facts as small, local, irreversible events — then nothing forces a global reconciliation, and the charge can survive. The programme has not shown which of these is true; it has shown that this is the whole question. There is a soft lean toward the local reading, because every building block in the framework is local and none lays down a global frame on its own — but that is a matter of which side carries the burden of proof, not a verdict.

7.2 Observability — Behind Occupancy

Observability is the second open question, and logically downstream of the first: it only arises if occupancy holds.

Suppose the closure sector is occupied. Does the surviving closure class survive the operational quotient?

N_hol → N_obs

Does the surviving holonomy become physically observable, or is it erased by commitment dynamics on the way to the observable sector? **(open, conditional on occupancy.)**

In plain terms. Suppose the channel *is* occupied — there is a real, irreducible closure memory. A second question follows: could anyone ever *detect* it? The framework has a filter — the operational quotient — that discards anything no measurement could distinguish. The surviving memory (N_hol) might be exactly what this filter throws away, so it never reaches the world of observable quantities (N_obs). So even a "yes" to occupancy would not guarantee a "yes" to observability. Two separate hurdles, and the framework is currently at neither — with occupancy the one that has to be cleared before the second even arises.

7.3 The Pre-Observable Register Seam

The register verdict of Section 6 is closed for the *observable* charge but leaves one honest seam, because both of its supporting floors are explicitly about measurement. The uniqueness and saturation theorems govern operational distinctions and the observable algebra; by construction they fall silent on a possible richer *pre-observable* substrate that no measurement reaches. A reader who wishes to keep a genuinely seven-valued charge alive has exactly one move: locate it in that pre-observable layer, below the measurement floor, where neither theorem reaches. **(open — but burden-assigned.)**

The cost of the move is exact. A seven-valued register relocated below the measurement floor is, by that relocation, no longer the observable closure holonomy the Gate-3 sector was defined to carry. It becomes a different object — one that shows up in no committed record. So the seam does not reopen the register of the *observable* charge, which stays $\mathbb{Z}_2$; it asks only whether a separate, unobservable seven-valued structure exists beneath it. The burden of exhibiting such a structure — and of accepting that it is not the Gate-3 holonomy — sits on the seven-valued reading.

In plain terms. The two-way-switch verdict is about what you can *measure*. The framework's foundations are careful to speak only about measurable things and stay silent about whatever the substrate might be doing *beneath* the level anything could ever detect. So there is exactly one way to keep a true seven-valued charge alive: put it down there, below the floor measurement reaches. The catch is that doing so changes what it is. The Gate-3 charge was, by definition, something you could in principle measure — that was the whole point. A seven-valued structure hidden below the measurement floor is no longer that thing; it leaves no fingerprint on any record. So the verdict for the charge the sector was actually about stays the two-way switch; the only open question is whether some separate, forever-hidden seven-valued structure sits underneath — and the burden of showing it is there falls on whoever wants to keep the seven.

8. The Current Frontier

The programme's frontier has changed significantly.

Early work focused on:

- probability,
- geometry,
- emergent geometric structure,
- gauge structure.

Current work focuses on:

- occupancy — specifically the Uniform Readout fork (UR-frame vs UR-comp),
- observability,
- the pre-observable register seam.

Several items that once sat on this frontier are discharged: completion is local-fact-generated (protected cycles persist into the continuum limit), conservation supplies flatness, the topology supplies availability, and — newly — the *register* of the charge is fixed: the observable holonomy is the two-valued class (Section 6), whatever its occupancy. What these do *not* settle is occupancy itself, which has turned out to rest on the readout rule rather than on protection, flatness, or register.

The central question is no longer:

Can geometry emerge?

Nor:

Can gauge structure emerge?

Nor:

Can gravity emerge?

The programme has largely addressed those.

The central question is now:

Does reading out the substrate consult one global, path-independent ledger — or only a path-dependent record of local comparisons?

This single property of the readout rule decides whether the substrate preserves any irreducible closure memory beyond local records. It is the final substrate-level residue, with observability waiting behind it — and, behind that, the pre-observable register seam. One thing is no longer in doubt: should a charge survive, its *register* is settled (the two-valued class, Section 6); what is open is whether it survives at all, and whether it could be seen.

In plain terms. It is worth appreciating how much the frontier has moved. Not long ago the live questions were the big foundational ones — can geometry, forces, gravity come out of this? Those have largely been answered within the framework. The frontier has since contracted to one sharp question, and that question is open: does reading the substrate mean consulting a single global ledger that every local comparison must agree with, or does it mean following a path-dependent trail of purely local comparisons that no global frame need reconcile? If the former, the closure memory vanishes; if the latter, it can survive. That one property of how records are read is now the whole of the occupancy question, with the question of whether any surviving memory could be *detected* (observability) waiting one step behind it.

---

9. Conclusion

The VERSF programme has evolved from a speculative emergence framework into a highly interconnected reconstruction programme.

A large number of apparent independent structures have been reduced to consequences of a comparatively small set of admissibility principles.

The framework now contains candidate derivations of:

- time,
- probability,
- Hilbert geometry,
- gauge structure,
- Lorentzian geometry,
- gravitational dynamics,
- particle classifications,
- closure transport.

Most major structural questions have been reduced to a small number of sharply defined substrate questions. Of the closure charge, two questions have been separated and answered to different depths.

The closure charge is a contextuality obstruction — a class in the first cohomology of the vacuum transport complex, nonzero exactly when the local pairwise comparisons fail to glue into a single global labelling. Its *register* — what the measurable charge is valued in — is now settled: the observable holonomy is the two-valued class, $\kappa \in H^1(\Gamma\_vac; \mathbb{Z}_2)$, forced by the uniqueness of the fold and the saturation of the observable algebra, and resting on the framework's most secure (Category I) tier. Its *occupancy* — whether the obstruction is nonzero at all ($A \subseteq B^{\perp}$) — remains open, reduced exactly to one property of the Uniform Readout rule: whether readout consults a single path-independent global ledger (UR-frame → $\kappa = 0$, sector empty) or only a path-dependent record of local comparisons (UR-comp → $\kappa \neq 0$ possible on a T1 topology). The seven keeps its genuine roles — the kernel of the transport group, and a count of constraints, loop channels, and gauge generators — but it is not the register of the observable charge.

The programme therefore now terminates at one open substrate question, with two waiting behind it:

1. Occupancy — does the Uniform Readout axiom assert UR-frame or UR-comp? Equivalently: does committing a fact ever lay down a global frame, or only a local comparison? **(open.)**
2. Observability — if occupied, does the surviving holonomy survive the operational quotient, N_hol → N_obs? **(open, conditional on occupancy.)**
3. The pre-observable register seam — whether a separate, unobservable seven-valued structure exists below the measurement floor. The observable charge is the two-valued class regardless; the burden of exhibiting any pre-observable seven-valued ontology, and of accepting that it is not the Gate-3 holonomy, sits on the seven-valued reading. **(open — burden-assigned.)**

A parsimony asymmetry has been recorded but not mistaken for a decision: on occupancy, every primitive in the generative chain is local, so Branch A (UR-frame) must exhibit a global-frame-fixing mechanism the local ontology does not supply, while Branch B (UR-comp) posits only its absence plus an available topology. This locates the burden on Branch A; it does not settle the question.

The next phase of the programme is therefore not broad expansion. It is closure — in the narrowed sense of fixing the one property of the readout rule on which occupancy turns, then settling observability, with the register of any survivor already known.

In plain terms. The programme set out to grow physics from a tiny seed, and it has reduced a sprawl of seemingly separate problems to a single connected structure. Time, probability, the geometry of quantum states, forces, the geometry of gravity, its dynamics, the classification of particles, and the transport of closure now sit inside one framework as candidate consequences of a short starting list — the gravity chain more mature than the matter chain, all of it resting on the admissibility assumptions. What is remarkable is how little is left, and how sharply it is stated. Two distinct questions about the closure charge have come apart. *What kind of charge would it be?* — answered: a two-way switch, not a seven-position dial, and this is one of the safest results in the framework, leaning only on its bedrock. *Is the charge there at all?* — still open, reduced to a single checkable property: when you read out the substrate, are you consulting one global ledger that every local comparison must match, or only a path-dependent trail of local comparisons? If the first, the channel is empty; if the second, it can carry a real, irreducible memory. There is a gentle lean toward the second, because every building block in the framework is local and none lays down a global frame by itself — but that is about where the burden of proof sits, not a verdict. Settle that one property, then ask whether any surviving memory could ever be seen, and whether anything seven-valued hides below the level measurement can reach. Those are the last doors — and the shape of whatever waits behind them is already known to be the simplest one.