# Matter from Persistent Fold Closure in VERSF

**Topological Commitment Defects, Interface Holonomy, and the Emergence of Particle Sectors**

*Keith Taylor — VERSF Theoretical Physics Programme*

---

## General reader summary

What is matter made of? Standard physics treats matter as stuff — particles like electrons and quarks — placed inside a passive container called space. This paper proposes something different. Matter is not stuff inside space. Matter and space both **emerge from** a deeper substrate — they are products of the same underlying dynamics, not arrangements within a pre-existing container.

Earlier papers in this programme built that substrate step by step. The starting point is the **Void**: not empty space, but the most primitive level of physical reality, where no distinctions yet exist. The first distinction is called a **Fold**. Folds can become irreversibly committed — locked in — through a topological trapping mechanism. Accumulated committed Folds form a **record field**, and from this record field, ordinary geometry and gravity emerge.

What was still missing was a substrate-level account of localized matter — why particles exist at all, why they have mass, charge, and spin, and why some of them behave as fermions (gaining a minus sign when swapped). This paper proposes a single answer:

*Matter **emerges from** the substrate as **Persistent Fold Defects** — stable topological structures formed by the substrate's irreversible commitment dynamics, configurations that once produced cannot be unmade.*

Ordinary disturbances in the substrate disperse and disappear. Persistent Fold Defects survive because their topological structure cannot be undone by any admissible local process. The same single object then does the work usually requiring six separate ingredients:

- it explains why particles **persist** — their topology is protected;
- it explains why particles carry **charge** — the knot has a non-trivial winding;
- it explains why particles have **mass** — the knot is stiff and costs energy to compress, move, or distinguish from its surroundings;
- it explains why some particles have **spin** — the knot's orientation behaves non-trivially when rotated;
- it explains **fermionic exchange** — swapping two such knots picks up a minus sign, the signature of fermions like electrons;

- it explains how matter **bends geometry** — the knot contributes to the record field that sources gravity.

What the paper does *not* yet do is give the dictionary between these knots and the specific particles physicists observe — electrons, quarks, neutrinos, and the rest. That dictionary is the next major step in the programme.

What the paper *does* establish is the kind of thing matter is: not a thing placed in space, but something that emerges from the same substrate dynamics that produce space itself. Matter and geometry are then no longer independent — they are two faces of the same substrate process.

---

# Contents



---

# Abstract

Earlier stages of the VERSF programme established the substrate architecture of physical reality: the Void as a zero-distinction constraint medium; the Fold as the minimal primitive distinction; the commitment interface as the physically active substrate layer; the formation of irreversible facts through topological trapping ($\beta_1 \geq 1$); the emergence of gauge transport from closure redundancy; the reconstruction of Lorentzian geometric structure from refinement-stable transport; and the emergence of Einstein-compatible gravitational dynamics from accumulated committed record structure.

One structural problem remained unresolved:

*Why does the framework contain stable localized matter sectors at all?*

The present paper addresses that problem.

We propose that matter corresponds to **Persistent Fold Defects (PFDs)** — localized regions of nontrivial commitment-interface closure topology whose transport holonomy, commitment density, and closure structure remain stable under admissible refinement flow. Ordinary fluctuations disperse under refinement. Persistent matter sectors survive because their homology class prevents continuous unwinding into the reversible substrate sector.

Six connected claims are developed:

1. **Persistent Fold Defects (§3).** Matter sectors correspond to localized closure complexes whose closure topology, holonomy, refinement-persistence, and stiffness are simultaneously nontrivial, and whose holonomy therefore cannot be globally trivialized.
2. **Topological Stability (§4).** PFDs are protected against local decay because non-contractible closure structure ($\beta_1 \geq 1$) cannot be removed by admissible local moves.
3. **Charge from Interface Holonomy (§5).** Conserved transport charge arises from stable closure holonomy around persistent defect sectors. U(1)-type charge appears as the continuum image of conserved closure circulation.
4. **Mass from Closure Stiffness (§6).** Mass emerges from four convergent contributions: commitment-density loading, closure-Hessian stiffness, confinement/localization cost, and persistent distinguishability content.
5. **Spinorial and Fermionic Structure (§7).** Spinorial transport around persistent defects generates orientation holonomy and exchange sign structure. The Finkelstein–Rubinstein mechanism forces antisymmetric exchange sectors for nontrivial persistent closure topology.
6. **Matter–Geometry Coupling (§8).** PFDs source the record field and therefore source emergent geometry. Stress-energy becomes the continuum coarse-grained image of persistent localized closure structure.

**Epistemic status.** The contribution is structural and ontological. The paper does *not* claim a derivation of the Standard Model spectrum, exact particle masses, scattering amplitudes, or experimentally complete matter phenomenology. It identifies the first unified substrate-level

object capable of simultaneously supporting persistence, mass, charge, spin, fermionicity, and geometric sourcing. Theorems 4.1 and 5.2 are proven from inherited primitives; Theorems 7.1, 7.2, and 8.1 are conditional on prior VERSF results; the unified mass expression of §6 is schematic and explicitly identified as such.

---

# 1. Introduction

The earlier VERSF papers established a progressively layered reconstruction of physical structure:

**Void → Fold → Fact → Record → Transport → Geometric structure → Gravity.**

From a common substrate architecture centred on the commitment interface, the programme has derived: finite distinguishability; irreversible commitment; topological trapping; closure transport; gauge redundancy; Maxwell admissibility; Lorentzian geometric emergence; Einstein-compatible dynamics; the Einstein–Hilbert action; and quantum kinematics (Hilbert space, Born rule, admissibility closure).

One layer remained absent: a **substrate-level ontology of localized matter sectors**.

Geometry had emerged. Transport had emerged. Gauge structure had emerged. Quantum kinematics had emerged. But the framework had not yet identified the substrate object that could simultaneously support:

- persistence (why particle-like sectors exist at all),
- mass (why they resist acceleration),
- charge (why they carry conserved transport quantities),
- spin (why some sectors are spinorial),
- fermionic exchange (why some sectors antisymmetrise),
- geometric sourcing (why matter curves geometry).

The present paper proposes that all six arise from a single substrate object: the **Persistent Fold Defect**.

A Persistent Fold Defect (PFD) is a localized closure complex on the commitment interface whose topology prevents refinement-driven dispersal into the reversible substrate sector. The central claim is:

*Matter is stable topological persistence on the commitment interface.*

Matter is not inserted into the emergent geometry as an independent primitive. It is a topologically protected excitation of the same Fold–Record substrate from which geometry itself emerges.

### Structural role in the programme

This paper occupies a specific architectural position. The substrate machinery *upstream* — Void, Fold, Fact, Record, Transport, Lorentzian geometric emergence, Gravity, and the quantum admissibility chain — has been established in earlier papers. The phenomenological content *downstream* — the explicit dictionary to Standard Model representations, the mass spectrum, scattering amplitudes, the SU(2)×U(1) and non-Abelian PFD sectors, the K = 7 closure quantum sector, and PFD second quantisation — sits in future companion papers.

The present paper is the bridge between the two. It is the first paper in the programme in which the substrate machinery directly produces the kind of object that downstream phenomenology is about: localized, persistent, charge-carrying, mass-carrying, spin-carrying excitations that source emergent geometry. Companion papers on the SU(2)×U(1) PFD sector, PFD second quantisation, and the PFD → Standard Model dictionary all sit *downstream* of this paper rather than parallel to it.

---

## 2. The Fold–Record Substrate (Inherited Structure)

The Fold itself, used throughout this paper as a substrate primitive, is not stipulated. It is established as the unique minimal primitive distinction by the no-alternative theorem of the companion *Uniqueness of the Minimal Distinction* paper: under finite distinguishability, finite encodability, internal closure, and background-independence, the only admissible primitive distinction (up to isomorphism) is the strongly connected two-cycle on two states — the Fold, with intrinsic $\mathbb{Z}_2$ automorphism. The substrate primitives we inherit below are therefore forced, not chosen, and the present paper builds on them under that status.

We inherit the following from earlier VERSF papers, used here as primitives without re-derivation:

- the Void $V$ as zero-distinction constraint medium;
- the commitment interface Σ as the physically active substrate layer;
- the fold variables $\phi_i = (\sigma_i, \omega_i)$ on the hexagonal K = 7 closure architecture, where $\sigma_i$ encodes distinguishability content and $\omega_i$ encodes orientation; these are the substrate image of the (b, d) basis of the per-fold Hilbert space $\mathscr{H}$_fold ≅ $\mathbb{C}^4$ established in the *One Fold* companion paper, with $\sigma_i \leftrightarrow b_i \in \{0, 1\}$ (the classical bit) and $\omega_i \leftrightarrow d_i \in \{+1, -1\}$ (the binary direction label forced by reversibility, Theorem D2 of that paper). The $\mathbb{Z}_2$ direction label that will appear in §7 as the source of spinorial structure is therefore already present at the substrate level, not introduced by dimensional emergence;
- the closure functional $C\ell$ on each closed plaquette ℓ;
- the irreversibility threshold C★ defining topological trapping;
- the committed record field ρ(x, t);
- the admissibility class $\mathcal{A}$ of permissible local refinement moves.

Closure transport along an oriented edge ⟨ij⟩ is generated by the local transport operator $U_{ij}$. The interface action used throughout is

$$S_{int} = \Sigma\_\langle ij\rangle\ \mathscr{L}(\phi_i, \phi_j) + \Sigma\_\ell\ C\ell + \Sigma\_\ell\ \Theta(C\ell - C\bigstar),$$

with $\mathscr{L}$ the local fold Lagrangian and Θ the irreversibility indicator. The third term implements topological trapping: once a plaquette exceeds C★, its closure becomes admissibility-fixed under all subsequent refinement.

Topological trapping occurs when

$$\beta_1(D) \geq 1$$

on a localized region $D \subset \Sigma$. The resulting committed folds contribute to ρ(x, t).

The question of the present paper is: under what additional conditions does such a localized closure structure remain *coherent, localized, and refinement-stable*, rather than dispersing as ordinary committed record?

---

# 3. Persistent Fold Defects

## Definition 3.1 — Persistent Fold Defect

A **Persistent Fold Defect** is a localized region $D \subset \Sigma$ satisfying all four of the following conditions:

**(P1) Nontrivial closure topology.** $\beta_1(D) \geq 1$, with at least one homology class $[\gamma] \in H_1(D)$ admissibility-fixed by topological trapping (i.e. every representative $\gamma \in [\gamma]$ bounds, in the one-sided sense, a 2-chain of trapped plaquettes ℓ with $C\ell \geq C\bigstar$; on Σ as a 2-manifold this is unambiguous).

**(P2) Nontrivial closure holonomy.** There exists at least one non-contractible loop $\gamma \in [\gamma]$ such that $U\gamma = \prod\_\langle ij\rangle\in\gamma\ U_{ij} \neq \mathbb{1}$.

**(P3) Refinement persistence.** Under the admissible refinement flow $R: \Sigma_n \rightarrow \Sigma_{n+1}$, the sequence $D_n = R^n(D)$ converges to a non-empty limit $D\infty$ *with* $\beta_1(D\infty) \geq 1$.

**(P4) Positive closure stability.** The second variation of $S_{int}$ restricted to localized deformations supported on $D$ is positive-definite: $\delta^2 S_{int}[D] \succ 0$.

## Interpretation

Ordinary commitment fluctuations fail (P1)–(P4) in identifiable ways:

- contractible commitment clusters fail (P1);
- closure-trivial loops fail (P2);
- diffusive committed regions fail (P3);
- saddle or unstable configurations fail (P4).

PFDs are the simultaneous solution set of all four. They are localized, topologically protected, holonomy-carrying, and energetically stable.

The structural claim of the paper is that *this is the ontological signature of matter*.

---

# 4. Topological Stability

### Lemma 4.0 — Admissibility-Fixed Homology *(proven)*

Let $\gamma$ be a 1-cycle in $D \subset \Sigma$ that bounds, in the one-sided sense, a 2-chain of trapped plaquettes ($C\ell \geq C\bigstar$) supported in $D$. Then $[\gamma] \in H_1(D)$ is fixed under every admissible chain map: for any admissible move $m \in \mathcal{A}$ inducing a chain map $m\_\#$ on cellular chains,

$m\_\#[\gamma] = [\gamma]$ in $H_1(D)$.

### Proof

An admissible move $m \in \mathcal{A}$ acts locally on the cellular structure of $\Sigma$. By the definition of $\mathcal{A}$ (§2), every such move preserves the irreversibility indicators $\Theta(C\ell - C\bigstar)$ on every plaquette $\ell$. In particular, no trapped plaquette is removed, retriangulated away, or merged into a non-trapped neighbour by an admissible move.

The induced action on cellular 1-chains is therefore restricted. Because admissible moves act *as the identity* on the sub-complex generated by trapped 2-cells and their boundaries, and only as a non-trivial chain map on the complement sub-complex of non-trapped 2-cells, any chain-level modification of $\gamma$ is supported in the complement sub-complex. The cellular cobordism implementing $\gamma \to m\_\#\gamma$ is therefore a chain of non-trapped 2-cells:

$m\_\#\gamma - \gamma = \partial c$, with $c$ a 2-chain of non-trapped 2-cells,

and so $[m\_\#\gamma] = [\gamma]$ in $H_1(D)$. ▢

### Remark 4.0.1

Lemma 4.0 is the bridge from the plaquette-level trapping established in §2 to the 1-cycle-level fixing required in Theorem 4.1. Without it, "admissibility-fixed homology class" would be a definition without a mechanism. With it, the trapping mechanism of §2 propagates from 2-cells to 1-cycles, and the $H_1$-level statement of Theorem 4.1 becomes proven rather than asserted.

## Theorem 4.1 — Persistent Topology Theorem *(proven, modulo §2 primitives and Lemma 4.0)*

Let $D \subset \Sigma$ be a localized closure complex satisfying (P1). Then $D$ cannot relax continuously into the reversible substrate sector $V$-sector through admissible refinement.

### Proof

Denote by $\mathcal{V}$ the reversible substrate sector, characterised by $\beta_1(\mathcal{V}) = 0$ on every localized neighbourhood. Suppose, for contradiction, that there exists an admissible refinement path

$D = D^{(0)} \to D^{(1)} \to \cdots \to D^{(N)} \subset \mathcal{V}$,

with each step $D^{(k)} \to D^{(k+1)}$ generated by a move in the admissibility class $\mathcal{A}$.

By definition of $\mathcal{A}$, each admissible local move preserves:

(i) edge-wise closure consistency, (ii) plaquette-wise irreversibility indicators $\Theta(C\ell - C\bigstar)$, (iii) the homology class of every admissibility-fixed loop (Lemma 4.0).

Condition (iii) is the operative one. By (P1), there exists a class $[\gamma] \in H_1(D)$ fixed by topological trapping (with $\gamma$ bounded on one side by trapped plaquettes, per Definition 3.1). Lemma 4.0 then gives

$[\gamma] \in H_1(D^{(k)})$ for all $k \leq N$.

In particular $[\gamma] \in H_1(D^{(N)})$. But $D^{(N)} \subset \mathcal{V}$ has $\beta_1 = 0$ on every localized neighbourhood, so $[\gamma] = 0$ — a contradiction.

Hence no such admissible path exists, and $D$ cannot relax continuously into the reversible sector. ▢

## Corollary 4.2 — Matter Persistence

PFDs survive arbitrarily many rounds of admissible refinement: topology prevents complete closure trivialisation, and refinement is required to preserve admissibility-fixed homology classes.

## Remark 4.3 — Non-perturbative content

The persistence statement is *non-perturbative*: it does not rely on the size of perturbations to $D$, nor on the magnitude of refinement steps. It is purely homological. Decay can only occur through *non-admissible* moves, which lie outside the substrate dynamics by construction.

---

# 5. Charge from Interface Holonomy

## 5.1 Setup

The transport companion paper established that admissible closure transport on Σ admits a continuum image given by a U(1) connection $A_\mu$ on the emergent geometric structure, with curvature $F_{\mu\nu}$ satisfying Maxwell admissibility. We use this image as the continuum description of closure transport around PFDs.

### Definition 5.1 — Closure Charge

The **closure charge** of a PFD $D$ is

$$Q_D = (1 / 2\pi) \oint_\gamma A_\mu \, dx^\mu,$$

where γ is a representative of the admissibility-fixed homology class in $H_1(D)$ and $A_\mu$ is the continuum closure-transport connection.

Independence of $Q_D$ from the choice of representative γ follows from gauge-equivalence of admissibility-preserving loop deformations within [γ].

### Theorem 5.2 — Charge Conservation *(proven, conditional on the U(1) transport image)*

Let γ be a spatial closure-loop representative enclosing $D$ entirely, and assume that no admissible closure current crosses the spatial boundary of γ over the interval $[t_1, t_2]$. If the closure-transport current $J^\mu$ satisfies the admissibility conservation law $\partial_\mu J^\mu = 0$, then

$$dQ_D / dt = 0$$

under admissible refinement evolution over $[t_1, t_2]$.

### Proof

Let $T(\gamma)$ denote the world-tube swept by γ under refinement evolution from $t_1$ to $t_2$. The boundary $\partial T(\gamma)$ consists of the two spatial loops $\gamma(t_1)$ and $\gamma(t_2)$ together with the timelike side $\partial_\| T(\gamma)$ traced by γ's spatial boundary.

**Step 1 — Bianchi.** By Stokes' theorem applied to $F$ on $T(\gamma)$,

$$\oint_{\gamma(t_2)} A - \oint_{\gamma(t_1)} A + \int_{\partial_\| T(\gamma)} F = \int_{T(\gamma)} dF = 0,$$

where the right-hand vanishing is the Bianchi identity $dF = 0$ on the emergent geometric structure, established in the Maxwell admissibility paper.

**Step 2 — No flux through the side.** By the spatial-enclosure assumption, no closure current crosses the spatial boundary of γ over $[t_1, t_2]$. Equivalently, the integrated flux through $\partial_\| T(\gamma)$ vanishes:

$$\int_{\partial_\| T(\gamma)} F = 0.$$

Combining Steps 1 and 2,

$$Q_D(t_2) - Q_D(t_1) = (1 / 2\pi) \left[ \oint_{\gamma(t_2)} A - \oint_{\gamma(t_1)} A \right] = 0.$$

**Step 3 — Identification with enclosed source.** Steps 1–2 establish that the holonomy $\oint_\gamma A$ is conserved as a winding quantity. *Identifying* this conserved holonomy as the conserved closure *charge* — i.e. as the enclosed source content — requires the source-side conservation law $\partial_\mu J^\mu = 0$. This guarantees that any admissible redistribution of source within $D$ over $[t_1, t_2]$ leaves the total source threaded by γ invariant, so that the geometric quantity $Q_D$ is also the physical charge.

The argument therefore has two distinct content layers: Steps 1–2 close the *geometric* side (the holonomy $\oint_\gamma A$ is conserved), and Step 3 closes the *identification* side (the conserved holonomy equals the conserved enclosed source). Both are needed: without Steps 1–2, $Q_D$ is not conserved; without Step 3, the conserved geometric quantity $\oint_\gamma A$ is not yet identified with the enclosed source content threaded by γ. ▢

## Remark 5.2.1 — Global vs local

The above is a global identity over the world-tube $T(\gamma)$. It becomes a *local* conservation law only after admissible coarse-graining over a family of nested world-tubes; the local form is the standard one and is not re-derived here.

## Remark 5.3 — Quantisation

Two independent facts together yield charge quantisation:

(a) **Integer-valued winding.** For any U(1) connection on a substrate with $H^1 \neq 0$, the closure-loop integral $(1 / 2\pi) \oint_\gamma A$ takes integer values on the admissibility-fixed homology class [γ]. This follows from the integer-valued homology pairing alone and does *not* require K = 7.

(b) **Fixing the unit.** The minimal nontrivial Wilson loop on the K = 7 hexagonal closure architecture fixes the unit of charge. Concretely, the *K = 7 Wilson Limit* (established in the Maxwell admissibility paper) identifies the smallest non-trivial Wilson loop on the hexagonal closure architecture as the unit closure quantum: every admissible closure circulation is an integer multiple of this minimal loop, and this integer multiple sets the charge unit for PFDs threading [γ].

(a) gives integer-valued $Q_D$; (b) fixes what an "integer" of charge actually is. A fully rigorous treatment of the K = 7 closure quantum sector and its relation to PFD topological invariants is deferred to a companion paper. *(Conditional / forward-referenced.)*

---

# 6. Mass from Closure Stiffness

Mass appears in four convergent forms across the VERSF programme. This section unifies them at the substrate level. **The unified mass expression (6.5) is explicitly schematic**: the four contributions are independently motivated, but their relative coefficients are not derived here.

## 6.1 Commitment density (ρ-loading)

A PFD generates localized committed-record density

$\rho_D(x, t) = \rho_{vac}(x, t) + \Delta\rho_D(x, t)$, with $\int \Delta\rho_D \, d^3x > 0$.

Elevated persistent commitment density resists further admissible compression: it is a mass-loading contribution to the substrate energetics.

## 6.2 Closure-Hessian stiffness

The second variation of the interface action restricted to $D$ defines the closure Hessian

$H_D = \delta^2 S_{int}[D]$.

By (P4), $H_D$ is positive-definite. Its lowest non-zero eigenvalue $\lambda_{min}(H_D)$ sets the inertial response of the defect to admissible perturbations and is the substrate-level analogue of inertial mass.

## 6.3 Confinement and localization cost

Localization requires suppression of closure dispersal. To leading order, the confinement cost of a defect localized on scale $L_D$ scales as

$E_{conf} \sim \hbar_{eff} / L_D$,

where $\hbar_{eff}$ is the substrate-emergent action scale established in the quantum admissibility papers. This is the substrate-level image of the standard uncertainty-principle confinement scale: closure-localized states pay an action cost set by $\hbar_{eff}$ and the inverse localization length, with $\hbar_{eff}$ fixed by the closure-quantisation results of the quantum admissibility chain and not introduced here as a free parameter. *(Substrate-level analogue of the uncertainty lower bound; the O(1) coefficient is set by the closure-quantisation results but not computed here.)*

### 6.4 Information content

A PFD stores persistent distinguishability structure. Its informational content scales with the number of admissibility-fixed folds,

$I\_D \sim N\_\text{fold}(D)$,

and contributes to the defect's energetic persistence via the Landauer-type bound established in the irreversibility papers.

### 6.5 Unified schematic mass expression

Combining the four contributions,

$m\_D \sim \alpha \cdot \int \Delta\rho\_D \, d^3x + \beta \cdot \lambda\_\text{min}(H\_D) + \gamma \cdot E\_\text{conf} + \delta \cdot I\_D$,

with coefficients α, β, γ, δ determined by substrate closure architecture.

**Epistemic status.** Expression (6.5) is *scaffolding*, not a derivation. The four terms are individually motivated by results established in earlier papers, but their relative weights, sign conventions, and cross-couplings are open. Moreover, the four contributions are *not in the same units* as written: $\int \Delta\rho\_D \, d^3x$ is a dimensionless count (or fold-density × volume), $\lambda\_\text{min}(H\_D)$ carries the dimensions of the second variation of $S_{\text{int}}$, $E\_\text{conf}$ is an energy, and $I\_D$ is dimensionless. The coefficients α, β, γ, δ are therefore *dimensional bridge constants*, not pure numerics; they carry most of the predictive content of (6.5). An honest derivation of $m\_D$ must produce these bridge constants from substrate primitives, not normalise them away. A genuine derivation requires:

- explicit construction of the canonical closure quantisation on *D*,
- identification of the Hessian zero-mode structure (translation, rotation, gauge),
- a controlled coarse-graining from the K = 7 architecture to the continuum mass operator,
- a substrate-level fixing of the dimensional bridge constants α, β, γ, δ,
- a demonstration that the four contributions are independent at the substrate level, or — failing that — an explicit accounting of their cross-couplings (in particular: $\lambda\_\text{min}(H\_D)$ and $E\_\text{conf}$ are both stiffness-like; ρ-loading and $N\_\text{fold}(D)$ both scale with committed-fold count on *D*, and these potential redundancies must be either ruled out or absorbed into the bridge constants; other redundancies among the four contributions cannot be excluded without a full derivation, since the four terms live on only three substrate axes — density, stiffness, information).

These are explicit programme targets and are not claimed here.

## Interpretation

Mass, at the substrate level, is *closure resistance*: the cost of compressing, displacing, or distinguishing persistent topological commitment.

## 7. Spinorial Structure and Fermionicity

The earlier spinorial programme established spinorial transport sectors, the Finkelstein–Rubinstein exchange mechanism, and partial CAR (canonical anticommutation relation) emergence. PFDs supply the missing substrate carrier for these structures.

### Theorem 7.1 — Spinorial Transport *(conditional on the orientation transport structure of §2 and the dimensional emergence chain)*

The argument below presumes the dependency chain

K = 7 closure architecture → emergent 3-space → SO(3) action on orientation bundle → $\mathbb{Z}_2$ double cover → spinorial sector,

established in the dimensional emergence paper. The 3-space rotation group SO(3) is *not* a primitive of the K = 7 substrate; it is the continuum image of orientation transport after dimensional emergence has acted. Theorem 7.1 inherits this dependency and is conditional on it.

The $\mathbb{Z}_2$ that appears in the double-cover step of this chain is *not created* by dimensional emergence. It is the continuum image, under dimensional emergence, of the substrate-level direction-label $\mathbb{Z}_2$ established at the cell level in the *One Fold* companion paper (Theorem D2: reversibility on one bit forces a $\mathbb{Z}_2$ direction label, realised on the interface as the ω-component of $\phi_i = (\sigma_i, \omega_i)$; see §2). Dimensional emergence does not introduce a new $\mathbb{Z}_2$; it *identifies* the substrate-level ω-$\mathbb{Z}_2$ with the $\pi_1(SO(3)) = \mathbb{Z}_2$ of the emergent 3-space rotation group. Whether this identification acts faithfully on each PFD is the content of Open question 7.1.1 below.

**Two orientation bundles.** The framework contains two distinct orientation bundles, related by dimensional emergence:

(i) the *substrate-level fold-orientation bundle* on Σ, with fibre $\omega_i$ as introduced in §2 — a K = 7-level primitive on the hexagonal closure architecture;

(ii) the *post-dimensional-emergence SO(3) orientation bundle* on the emergent 3-space — the continuum image of (i) after dimensional emergence acts.

Bundle (ii) is the continuum image of bundle (i) under the dimensional-emergence map: (i) persists at the substrate level; (ii) is its coarse-grained continuum image and is what carries the SO(3) structure group. The SO(3) structure group is *not* present at the substrate level. The statement U_{2π} = −𝟙 below is therefore a *continuum-image* statement on bundle (ii), and its lift to bundle (i) is governed by the dimensional emergence map. A substrate-level statement directly on $\omega_i$ is a stronger claim than is made here; it requires the explicit dimensional-emergence dictionary and is deferred to the companion paper.

Let $D$ be a PFD whose admissibility-fixed homology class $[\gamma]$ generates orientation holonomy $U\gamma$ acting nontrivially on the (post-dimensional-emergence) orientation bundle over $\Sigma$. If this bundle admits a non-trivial double cover under closed-loop transport, then transport around $2\pi$-rotations of $D$ satisfies

$$U_{2\pi} = -\mathbb{1}$$

on the spinorial sector, and the transport representation is spinorial rather than tensorial.

## Proof sketch

Closed orientation transport defines a representation of $\pi_1(SO(3))$ (or its higher-dimensional analogue) on the orientation fibre over $D$. The double-cover structure of SO(3) by Spin(3) ensures that representations split into tensorial (factoring through $\pi_1 = \mathbb{Z}_2$ trivially) and spinorial (faithful on $\mathbb{Z}_2$) sectors. PFDs whose orientation transport is faithful on the $\mathbb{Z}_2$ factor land in the spinorial sector and satisfy $U_{2\pi} = -\mathbb{1}$. ▢

## Open question 7.1.1

Whether the K = 7 orientation transport bundle, after dimensional emergence, is faithfully a Spin(3) bundle — or whether it is a more general (possibly higher-spin or non-faithful) representation of the local orientation group — is not established here. Theorem 7.1 invokes the splitting into tensorial and spinorial sectors; the *existence* of both sectors on K = 7 is a separate question, deferred to the dimensional emergence companion paper.

## Theorem 7.2 — Fermionic Exchange *(conditional on Theorem 7.1 and on the FR exchange mechanism)*

Two spinorial PFDs of the same closure type satisfy antisymmetric exchange:

$$U_{exch} = -\mathbb{1}$$

on their joint orientation sector.

## Proof sketch

Exchange of two PFDs in the emergent 4-manifold is implemented by braiding their world-tubes within a fixed spatial slice. The Finkelstein–Rubinstein construction supplies the operative homotopy: in emergent 3-space, the exchange path that interchanges two PFDs while leaving their orientation frames coherently parallel-transported is *homotopic* to a path in which one PFD remains fixed while the other undergoes a $2\pi$ rotation of its orientation frame. This homotopy is the substrate-level content of the FR identification and is what carries the sign.

Applying Theorem 7.1 requires moving the holonomy along the FR homotopy from the exchange path to the $2\pi$-rotation path. The licensing fact is that the holonomy in the exchange-sign sector is $\mathbb{Z}_2$-valued — $U_{exch} \in \{\pm\mathbb{1}\}$ — so homotopy invariance is automatic in this sector: any curvature-flux correction through the FR interpolating surface lives in a continuous group and therefore can only contribute trivially to a $\mathbb{Z}_2$-valued holonomy. The exchange path and the 2π-rotation path therefore carry the same sign. Applying Theorem 7.1 to the 2π-rotation representative of this homotopy class then gives $U_{2\pi} = -\mathbb{1}$ on the spinorial sector, and

$$U_{exch} = U_{2\pi} = -\mathbb{1}$$

on the spinorial sector. ▢

### Remark 7.3 — Spin–statistics at the substrate level

Theorems 7.1–7.2 together implement spin–statistics at the substrate level: spinorial PFDs are necessarily antisymmetric under exchange, and tensorial PFDs are necessarily symmetric. The connection between rotation topology and exchange topology is *intrinsic* to the Fold–Record substrate; it is not imposed by hand.

A full reconstruction of CAR algebras and Fock structure on the space of multi-PFD configurations is deferred to a companion paper on PFD second quantisation.

---

# 8. Matter–Geometry Coupling

## 8.1 Sourcing of the Record Field

PFDs are persistent contributors to the committed record field. For PFD populations that are *either* non-overlapping in support *or* coarse-grained over admissibility scales larger than the maximum PFD localization length $\max_\alpha L_{D_\alpha}$, the record field decomposes additively:

$$\rho(x, t) = \rho_{vac}(x, t) + \Sigma_\alpha \rho_{D_\alpha}(x, t),$$

where the sum runs over all PFDs $D_\alpha$ *present in the region. Each* $\rho_{D_\alpha}$ is localized, refinement-stable, and positive on its support. Additivity in (8.1) is *not* assumed a priori on overlapping supports. At coarse-graining scales below $\max_\alpha L_{D_\alpha}$*, overlap corrections enter the residue and can be non-negligible; (8.1) holds in the coarse-graining limit at scales above* $\max_\alpha L_{D_\alpha}$, not in the strict pointwise sense at sub-localization scales. A treatment of sub-localization overlap effects is deferred to the multi-PFD companion paper.

## 8.2 Sourcing of Emergent Geometry

The gravity programme established the fundamental rank-2 commitment field $\Phi_{\mu\nu}$ as the primary geometric primitive, with the emergent metric given by

$$g_{\mu\nu} = (1 / \lambda\bigstar)\, \Phi_{\mu\nu}$$

(Theorem 10.0 of the gravity programme), where $\lambda\bigstar$ is the substrate stiffness scale fixing the gravitational coupling.

The scalar record density $\rho$ of §6.1 is *not* the more fundamental object. Two complementary relationships connect $\rho$ and $\Phi_{\mu\nu}$, operating at different scales:

**Macroscopic trace projection:** $\rho = g^{\mu\nu} \Phi_{\mu\nu}$.

**Microscopic kernel reconstruction:** $\Phi_{\mu\nu}(x) = \int K_{\mu\nu}(x, x') \rho(x', t') d^4x'$,

with $K_{\mu\nu}$ the conservation-fixed minimal covariant kernel of the gravity programme — built from the metric, covariant derivatives, and a single massive Green function at leading derivative order, with conservation $\nabla^\mu \Phi_{\mu\nu} = 0$ (away from sources) fixing its internal coefficient structure. The trace projection is not a stipulation: it is the gravity programme's identification of the macroscopic record density of §6.1 with the metric trace of the fundamental commitment field, established together with the rank-2 selection theorem and the metric definition $g_{\mu\nu} = (1/\lambda\bigstar) \Phi_{\mu\nu}$ via closure-consistent parallel transport (gravity programme, Theorem 10.0 and surrounding development). The fundamental geometric object is therefore $\Phi_{\mu\nu}$; $\rho$ is its macroscopic trace; and $K_{\mu\nu}$ bridges them at the microscopic level.

## Theorem 8.1 — Stress-Energy Emergence *(conditional on the gravity programme)*

A population of PFDs contributes additively to the fundamental commitment field $\Phi_{\mu\nu}$ through the closure-response kernel acting on their localized record contributions:

$$\Phi_{\mu\nu}^{(\mathrm{PFDs})}(x) = \int K_{\mu\nu}(x, x') \left[ \Sigma_\alpha \rho_{\{D_\alpha\}}(x', t') \right] d^4x',$$

with $K_{\mu\nu}$ the conservation-fixed closure-response kernel of §8.2. The continuum stress-energy contribution of these PFDs is the gravity programme's stress-energy functional $T^{(\Phi)}_{\mu\nu}$ — *quadratic* in $\Phi$ — evaluated on the PFD-sourced sector:

$$T_{\mu\nu}^{(\mathrm{matter})} = T^{(\Phi)}_{\mu\nu}[\Phi^{(\mathrm{PFDs})}].$$

This stress-energy enters the Einstein equation of the gravity programme,

$$G_{\mu\nu} + \Lambda g_{\mu\nu} = \kappa T^{(\Phi)}_{\mu\nu},$$

as the matter contribution sourced by Persistent Fold Defects.

## Proof sketch

Each PFD generates a positive, localized, refinement-stable record contribution $\Delta\rho_{\{D_\alpha\}}$ on its support (§6.1, §8.1). The map $\rho \to \Phi_{\mu\nu}$ is supplied by the closure-response kernel $K_{\mu\nu}$ of the gravity programme: the unique conservation-fixed minimal covariant kernel built from the metric, covariant derivatives, and a single massive Green function at leading derivative order (§8.2). Linearity of the convolution against $K_{\mu\nu}$ gives the additive $\Phi_{\mu\nu}^{(\mathrm{PFDs})}$ above. The gravity programme then constructs $T^{(\Phi)}_{\mu\nu}$ as the quadratic-in-$\Phi$ functional that sources Einstein curvature. Lovelock's theorem in this construction fixes the *metric-derivative* content of the gravitational sector — its unique divergence-free, symmetric, low-derivative form in 4D; the *$\Phi$-dependence* of $T^{(\Phi)}_{\mu\nu}$ is fixed independently by conservation and by the substrate-level closure-quantisation structure of the gravity programme. Evaluating this functional on $\Phi^{(PFDs)}$ yields the PFD stress-energy contribution.

Two objects must be kept distinct throughout: $\Phi_{\mu\nu}$ (linear in record density, metric-proportional, fundamental) and $T^{(\Phi)}_{\mu\nu}$ (quadratic in $\Phi$, source of Einstein curvature). The gravity programme also keeps a third object distinct — the closure current $\mathcal{C}^{\mu\nu}$ (linear in $\Phi$, leading-order record transport) — which does not enter the present paper and whose definition lives upstream. Conflating $\Phi_{\mu\nu}$ with $T^{(\Phi)}_{\mu\nu}$ — for instance by treating $\delta\Phi_{\mu\nu} / \delta g^{\mu\nu}$ as a stress-energy expression — collapses the construction, since $\Phi_{\mu\nu}$ is proportional to the metric and $\delta\Phi_{\mu\nu} / \delta g^{\mu\nu}$ therefore captures the geometric proportionality rather than the physical content of the source. ▢

### Remark 8.2 — Energy conditions

Positivity of $\Delta\rho_{D_\alpha}$ on its support, combined with positive-definiteness of $H_{D_\alpha}$ (condition P4), produces a positive contribution to the energy-density component $T_{\mu\nu}^{(matter)} u^\mu u^\nu$ *of the coarse-grained stress-energy tensor along any timelike* $u^\nu$ *in the emergent geometric structure. The mechanism is two-step. First, positive* $\Delta\rho_{D_\alpha}$ produces a positive contribution to $\Phi_{\mu\nu}^{(PFDs)}$ through the conservation-fixed closure-response kernel $K_{\mu\nu}$, which carries the appropriate sign on the trace channel (the $g_{\mu\nu} G_m$ term of the kernel ansatz). Second, the quadratic-in-$\Phi$ stress-energy functional $T^{(\Phi)}_{\mu\nu}$ of the gravity programme inherits this positivity on the timelike eigendirection, since quadratic forms of positive symmetric arguments retain positivity on aligned projections. The resulting weak-energy-condition analogue is therefore a *consequence* of (P4) and the $\rho$-loading positivity of §6.1, mediated by the gravity programme's construction of $\Phi_{\mu\nu}$ and $T^{(\Phi)}_{\mu\nu}$. A full treatment of all standard energy conditions (null, strong, dominant) for PFD-sourced stress-energy is open. *(Conjectural / programme target.)*

---

# 9. Interpretation

The framework now possesses a unified substrate ontology:

**Void → Fold → Fact → Persistent Closure → Matter → Geometric structure → Gravity.**

Matter is not an independently inserted primitive of the theory. It is a *protected mode* of the same Fold–Record substrate from which all other structure emerges:

- closure topology gives **persistence** *(proven, §4)*,
- closure holonomy gives **charge** *(proven conditional on the U(1) transport image, §5)*,
- closure stiffness gives **mass** *(schematic, §6)*,

- orientation transport gives **spin** *(conditional on dimensional emergence, §7)*,
- braiding holonomy gives **fermionic exchange** *(conditional on Theorem 7.1 and the FR mechanism, §7.2)*,
- record sourcing gives **geometric coupling** *(conditional on the gravity programme, §8)*.

These six entries are *not at equal epistemic standing* — only the first is proven from inherited primitives without further conditioning; the others rest on results imported from elsewhere in the programme or are explicitly schematic. They are nonetheless six images of one substrate object — the Persistent Fold Defect — viewed through six different admissible coarse-grainings.

---

# 10. Open Problems

The present paper does not yet derive:

- the explicit map from PFD topological invariants to Standard Model representations,
- the precise relationship between PFD topological invariants and the K = 7 closure quantum sector structure (the bridge object that connects this paper to charge quantisation, §5.3),
- the precise mass spectrum of Standard Model particles,
- non-Abelian closure sectors (SU(2), SU(3)) carried by PFDs of higher topological type,
- scattering amplitudes between PFDs,
- the cross-term structure of multi-PFD stress-energy: because $\Phi\_\mu\nu\^(PFDs) = \Sigma\_\alpha\ \Phi\_\mu\nu\^(D\_\alpha)$ is additive but $T\^(\Phi)\_\mu\nu$ is quadratic in Φ, the stress-energy of a multi-PFD population carries cross terms between distinct PFDs, and whether these reproduce standard matter–matter gravitational interaction in the appropriate limit is open,
- the renormalization-group flow of PFD sectors under continuum coarse-graining,
- confinement dynamics in the SU(3) closure sector,
- the precise coefficients α, β, γ, δ of the schematic mass expression (6.5),
- a full second-quantised Fock structure on the space of multi-PFD configurations.

These are explicit programme targets. A companion paper will address the SU(2)×U(1) PFD sector; a further companion paper will address PFD second quantisation and the emergent CAR algebra.

**The defining open problem.** Of these, the *first* item — the explicit map from PFD topological invariants to Standard Model representations — is the single largest open question for the next stage of the programme. The paper supplies the substrate ontology of matter; it does not yet supply the dictionary between PFD topological invariants and observed particle content.

Upstream work now sharpens this problem to a precise form. The *One Fold* companion paper establishes $\mathscr{H}$_fold ≅ $\mathbb{C}^4$ with basis |b, d⟩ (Theorem T1), the $\mathbb{C}^4 = \mathbb{C}^1 \oplus \mathbb{C}^3$ decomposition (Lemma GG2, from T1 + V1), and SU(3) × SU(2) × U(1) emerging as the commutant of the 3⊕1-block hopping matrix K acting on $\mathbb{C}^4$ (Appendix D.5). The complex, irreducible, non-self-conjugate representation property of the $\mathbb{C}^3$ sector is upgraded in *Companion Paper III* from a free axiom to a constrained selection — the unique representation type compatible with the substrate-level $\mathbb{Z}_2$ direction label and minimal distinguishability.

The PFD → Standard Model dictionary therefore takes the precise form:

*Identify which PFD topological invariants correspond to which orbit types of the SU(3) × SU(2) × U(1) action on $\mathscr{H}$_fold = $\mathbb{C}^4$ at the substrate cells supporting the defect.*

This is not the full dictionary, but it specifies the *shape* of the dictionary — and reduces the open problem from "find some correspondence between PFDs and observed particles" to "classify PFD topological types by their action on the substrate-level $\mathbb{C}^4$ structure." That is the bridge between the substrate-level claim *matter is persistent fold closure* and the phenomenological claim *matter is the Standard Model spectrum*, and constructing it is the defining target of the next phase.

**PFD gravitational self-energy and substrate renormalization.** Because *T*^(Φ)_μν is quadratic in Φ, a population of PFDs generates two distinct classes of contribution to the gravitational sector: *cross terms* between distinct PFD sectors (Φ^(D_α) · Φ^(D_β) with α ≠ β) and *diagonal self-interaction terms* (Φ^(D_α) · Φ^(D_α)). The cross terms govern multi-PFD gravitational interaction — Newtonian-limit recovery and equivalence-principle behaviour — and are listed as the cross-term item in the §10 bullet list above. The diagonal terms constitute *PFD gravitational self-energy*: the contribution of a single PFD to its own gravitational sourcing.

Whether these self-energy contributions remain finite under admissible coarse-graining, require substrate-level renormalization, or are automatically regularized by the finite closure scale of the K = 7 substrate is open. Three specific sub-questions:

- the presence or absence of ultraviolet divergence in the PFD self-sector under admissible refinement,
- the role of the localization scale *L*_D (§6.3) in self-energy regulation, and
- the relationship between closure-Hessian stiffness (§6.2, the substrate analogue of inertial mass) and gravitational mass renormalization on a PFD — i.e. whether the substrate ontology preserves equality of renormalized inertial and gravitational mass.

This issue is structurally new to the quadratic-in-Φ reconstruction of §8 — it does not arise in linear-in-Φ ansätze — and is the first genuinely nonlinear field-theoretic consequence of the substrate matter ontology developed here. The framework now confronts self-energy, interaction terms, equivalence-principle structure, and renormalization as standard issues that appear once a substrate ontology behaves as an interacting field theory, and these are now explicit programme targets rather than implicit ones.

**Methodological note on the §6.5 bridge constants.** The upstream *Area Scaling* companion paper (technical appendix to Companion Paper III) establishes a methodological precedent directly relevant to §6.5. There, the cosmological vacuum-energy parameter f ∼ ℓ_P / R is derived from admissibility primitives alone — locality of the Hamiltonian, finite on-site dimension ($\mathbb{C}^4$), TPB coarse-graining, and three spatial dimensions — without invoking either a full canonical quantisation or gravitational input. The *exponent* is forced by structure (area-to-volume scaling in d = 3); only the O(1) prefactors require further work.

This precedent suggests a partial route to the §6.5 dimensional bridge constants α, β, γ, δ. The four mass contributions (ρ-loading, closure-Hessian stiffness, confinement, information content) live on three substrate axes (density, stiffness, information). Scaling arguments using admissibility primitives plus the area/volume structure of the K = 7 substrate may resolve their *relative* magnitudes and the *scaling exponents* of their absolute scales in advance of the canonical closure quantisation listed as a prerequisite in §6.5; the O(1) prefactors remain open and require the canonical quantisation. This does not eliminate the prerequisites listed there; it identifies an intermediate route to the bridge constants that does not require closing all of them simultaneously. Whether the four contributions admit a scaling derivation of this kind is itself a concrete programme target.

---

# 11. Conclusion

Earlier VERSF papers reconstructed irreversible fact formation, gauge transport, Lorentzian geometric emergence, Einstein-compatible gravity, and quantum admissibility structure. The present paper adds the missing ontological layer: **localized matter sectors**.

The central result is structural:

*Matter corresponds to persistent topological closure defects of the commitment interface.*

Persistent Fold Defects survive refinement; carry closure holonomy; source gauge transport; generate mass through closure stiffness; acquire spinorial exchange structure under orientation transport; and source emergent geometry through the record field.

The resulting architecture closes the substrate genealogy:

**Void → Fold → Fact → Persistent Closure → Matter → Geometric structure → Gravity.**

Matter and geometry now emerge from the same substrate process — the irreversible stabilisation of distinguishability on the commitment interface. The framework no longer requires matter as an independent input. It is a topologically protected output of the same dynamics that produces everything else.