

The Void, The Fold, and the Derivation of Spacetime, Quantum Correlations, and the Standard Model Gauge Group

VERSF Theoretical Physics Program

For the General Reader

Physics has a long-standing puzzle at its foundations: where do the rules come from? Most theories begin by assuming a stage — space, time, a quantum field — and then describe what happens on it. But that leaves an awkward prior question unanswered: why does the stage have the properties it does? Why is time irreversible? Why do physical facts, once established, stay established? Why, if the universe has some discrete microscopic structure, can we not detect a preferred direction in space by travelling fast enough? And why do particles that have interacted remain correlated — even when separated by vast distances — in ways that no ordinary hidden mechanism can explain?

This paper addresses all of these questions from a single starting point.

The starting point is the void — not empty space, but something more fundamental: a substrate that is incapable of sustaining any distinction whatsoever. No difference can persist in it. No fact can be recorded in it. It is defined entirely by what it cannot represent.

If the void cannot contain distinctions, a question immediately arises: what happens to a distinction when it forms? It cannot stay. It must go somewhere. The answer this framework proposes is that it goes to a boundary — a structural edge between the void and a region where distinctions *can* persist. That boundary is called the fold.

The fold is the structural boundary that Axiom V1 requires. Physical structure accumulates at the fold boundary — specifically, the distinguishable content that cannot be represented within V. This paper derives consequences of that structure.

Once a distinction is deposited at the fold boundary, it becomes a *fact* — an irreversible record of something that has happened. Facts build up. Their accumulation is what we experience as time. Their structure is what we experience as space. The specific way facts can and cannot relate to one another — captured by a mathematical object called a partial order — turns out to be exactly the kind of structure that produces the geometry of spacetime as we know it, including the rule that no measurement can detect a preferred direction or a universal "now."

The quantum correlation puzzle — why entangled particles behave as if coordinated regardless of the distance between them — resolves in a similar way. Before any measurement occurs, an

entangled pair is not two separate particles each carrying a hidden answer. It is a single unresolved relation in the void substrate — one distributed object, not two local ones. When measurements happen at each end, both outcomes are produced by a single shared geometric constraint. This is why the correlations are stronger than any classical hidden scheme allows, but not stronger than quantum mechanics predicts: the geometry of the constraint imposes an exact upper limit — the Tsirelson bound.

The deepest result in this paper is that these two features — the structure of spacetime and the statistical ceiling on quantum correlations — are not independent. Both arise from a single principle about the fold: that observers can only ever access the fold through local projections, and that the rules governing those projections are the same in both cases. When you ask "why can't I detect a preferred direction in space?", you are asking the same question as "why can't entangled particles be more correlated than quantum mechanics allows?" Both answers are: because the fold only exposes invariant geometric overlaps, and conceals everything else. This is the Fold Projection Symmetry Principle — one of the central results of the paper.

There is a third result that may be the most striking of all. The Standard Model of particle physics — the theory that describes every known elementary particle and the forces between them — is governed by a specific mathematical symmetry group called $SU(3) \times SU(2) \times U(1)$. In every existing approach to fundamental physics, this group is either assumed outright or selected from an enormous landscape of possibilities using additional input. No existing framework derives it directly from the geometry of a pre-geometric boundary without additional structural assumptions.

VERSF derives it from the fold boundary geometry. The reasoning is as follows. The fold creates a two-dimensional boundary surface embedded in a three-dimensional emerging space. At any point on that boundary, there are exactly three kinds of local geometric information: how far apart two nearby events are (a single number), which way the boundary is oriented (a two-component spinorial object), and how the boundary curves (a three-component shape description). These three channels require complex vector spaces of dimensions 1, 2, and 3 respectively. A mathematical theorem then uniquely identifies the smallest symmetry groups that can act on those spaces: $U(1)$, $SU(2)$, and $SU(3)$. Their combination is exactly $SU(3) \times SU(2) \times U(1)$ — the Standard Model gauge group. No gauge structure is assumed at any point. It is what the local geometry of fact formation requires.

What remains to be proved mathematically is stated precisely. The paper is honest about the gap between the conceptual argument — which is complete — and the formal derivations — which constitute the program for the next stage of research. These results follow under specific structural assumptions described in the technical sections, several of which remain to be formally derived.

Table of Contents

1. The Problem
2. The Void
 - 2.1 Definition
 - 2.2 What the Void Is Not
 - 2.3 Reversibility as a Consequence
 - 2.4 Topological Characterisation
3. The Fold
 - 3.1 Symmetry Cannot Contain Its Own Violation
 - 3.2 The Fold as Structural Boundary
 - 3.3 The Inversion
4. The Topology of Facts
 - 4.1 What a Fact Requires
 - 4.2 Why the Fold Creates Facts
 - 4.3 Dimensional Signature
5. From Facts to Causal Structure
 - 5.1 Commitment Events
 - 5.2 The Commitment Structure
 - 5.3 Fold Admissibility
 - 5.4 The Partial Order Is Not a Foliation
6. The Lorentz Objection and Its Hidden Assumption
 - 6.1 The Standard Argument
 - 6.2 The Hidden Assumption
7. How Lorentz Invariance Emerges
 - 7.1 The Interval from Commitment Counting
 - 7.2 Light Cones from Upper-Set Boundaries
 - 7.3 Time Dilation from Chain Length
 - 7.4 No Observable Preferred Frame
8. Relation to Causal Set Theory
9. Non-Factorizable Commitment and Bell Non-Locality
 - 9.1 Why Local Bit Assignment Fails
 - 9.2 The VERSF Alternative
 - 9.3 Settings as Commitment Directions
10. The Tsirelson Bound from Commitment Geometry
 - 10.1 Five Constraints on the Void Substrate
 - 10.2 Deriving the Singlet Correlation Law
 - 10.3 The Tsirelson Bound
 - 10.4 Why Superquantum Correlations Are Excluded
 - 10.5 The Residual Gap
11. Fold Projection Symmetry: The Common Root
 - 11.1 The Shared Structure of Both Problems
 - 11.2 The Fold Projection Symmetry Principle
 - 11.3 The Spacetime Sector: Lorentz Covariance
 - 11.4 The Entanglement Sector: The Tsirelson Bound
 - 11.5 The Unified Fold Symmetry Conjecture

- 11.6 What a Full Proof Would Require
 - 12. The Standard Model Gauge Group from Fold Boundary Geometry
 - 12.1 The Prior State of the Derivation
 - 12.2 The Three Geometric Channels
 - 12.3 Channel 1: Scalar Separation $\rightarrow U(1)$
 - 12.4 Channel 2: Spinorial Orientation $\rightarrow SU(2)$
 - 12.5 Channel 3: Extrinsic Curvature $\rightarrow SU(3)$
 - 12.6 Why the Boundary Cannot Be Minimal ($H \neq 0$)
 - 12.7 Exhaustiveness and the Full Derivation
 - 12.8 Honest Assessment
 - 13. The Synthesis
 - 14. Addressing Standard Objections
 - 15. Open Problems
 - 16. Conclusion
-

Abstract

Three foundational VERSF papers — on the void substrate, the one-fold origin of structure, and the topological threshold for fact formation — converge on a unified account of how physical structure emerges from a zero-entropy substrate. This synthesis derives four interconnected consequences. (1) *Lorentz covariance*: the VERSF commitment structure admits no global foliation, so no preferred frame is available to any observer; the standard discrete-spacetime objection is dissolved. (2) *Tsirelson bound*: five structural constraints on joint commitment amplitudes select the quantum singlet correlation law as the minimal compatible solution and yield $|S| \leq 2\sqrt{2}$. (3) *Unified Fold Symmetry Conjecture*: both results are organised by three projection principles (FPS1–FPS3) governing how observable commitment statistics depend only on invariant quadratic projections of the fold; this conjecturally unifies Lorentz covariance and the Tsirelson bound as two representations of a single fold symmetry group G_F . (4) *Standard Model gauge group*: by the fundamental theorem of surface theory, the 2-dimensional fold boundary embedded in 3-dimensional space decomposes exhaustively into three geometric channels — scalar separation (C^1), spinorial orientation (C^2), and extrinsic curvature (C^3); the minimum compact connected Lie groups acting faithfully on these spaces are $U(1)$, $SU(2)$, $SU(3)$ (Theorem VIII), giving $SU(3) \times SU(2) \times U(1)$. Three limitations are stated explicitly: the $H \neq 0$ condition (required for C^3) is structurally motivated but not yet formally derived; the result is $SU(2)$, not $SU(2)_L$, as chirality remains to be established; and the relationship to the BCB framework's independent information-theoretic derivation of the same gauge group requires comparative analysis. A further foundational limitation is that the complex inner-product structure used in both the Tsirelson and gauge-channel derivations is currently inherited from the companion fold-realization analysis rather than derived directly from Axiom V1, and therefore the Tsirelson and gauge-channel derivations depend on this assumption.

Keywords: void substrate, fold topology, topological facts, causal partial order, Lorentz invariance, Tsirelson bound, Bell inequality, fold projection symmetry, non-factorizable

commitment, Standard Model gauge group, $SU(3)\times SU(2)\times U(1)$, extrinsic curvature channels, pre-geometric unification, BCB framework

Scope of this paper. The present work is a synthesis paper within the VERSF research program. Its purpose is not to present a completed unified theory but to show how several previously developed structural results follow from a single pre-geometric starting point. The paper derives three classes of consequences from the void–fold framework: (i) the causal partial-order structure leading to Lorentz covariance under coarse-graining, (ii) the geometric constraints on joint commitment amplitudes leading to the Tsirelson bound, and (iii) the geometric channel decomposition of the fold boundary leading to the gauge group $SU(3)\times SU(2)\times U(1)$. In each case the logical structure is established and the remaining derivations are identified explicitly as open problems.

1. The Problem

Most foundational physics frameworks take their substrate as a brute given — the quantum vacuum, spacetime, a Hilbert space — and derive structure from it. This is legitimate but leaves open a prior question: why does the substrate have the character it does, and how do irreversibility, time, and physical law arise from it?

VERSF addresses this by proposing that the substrate is not an arbitrary starting point but a constrained one: a zero-entropy void that cannot tolerate asymmetry. The question then becomes not *why does the void exist?* — that is a metaphysical question no physical framework can answer — but rather: given a zero-entropy substrate defined by its intolerance of asymmetry, what necessarily follows?

The answer developed across the VERSF program is that several key structures of modern physics follow from the same pre-geometric starting point: the character of spacetime, the absence of any preferred inertial frame, and the precise statistical limits on quantum correlations.

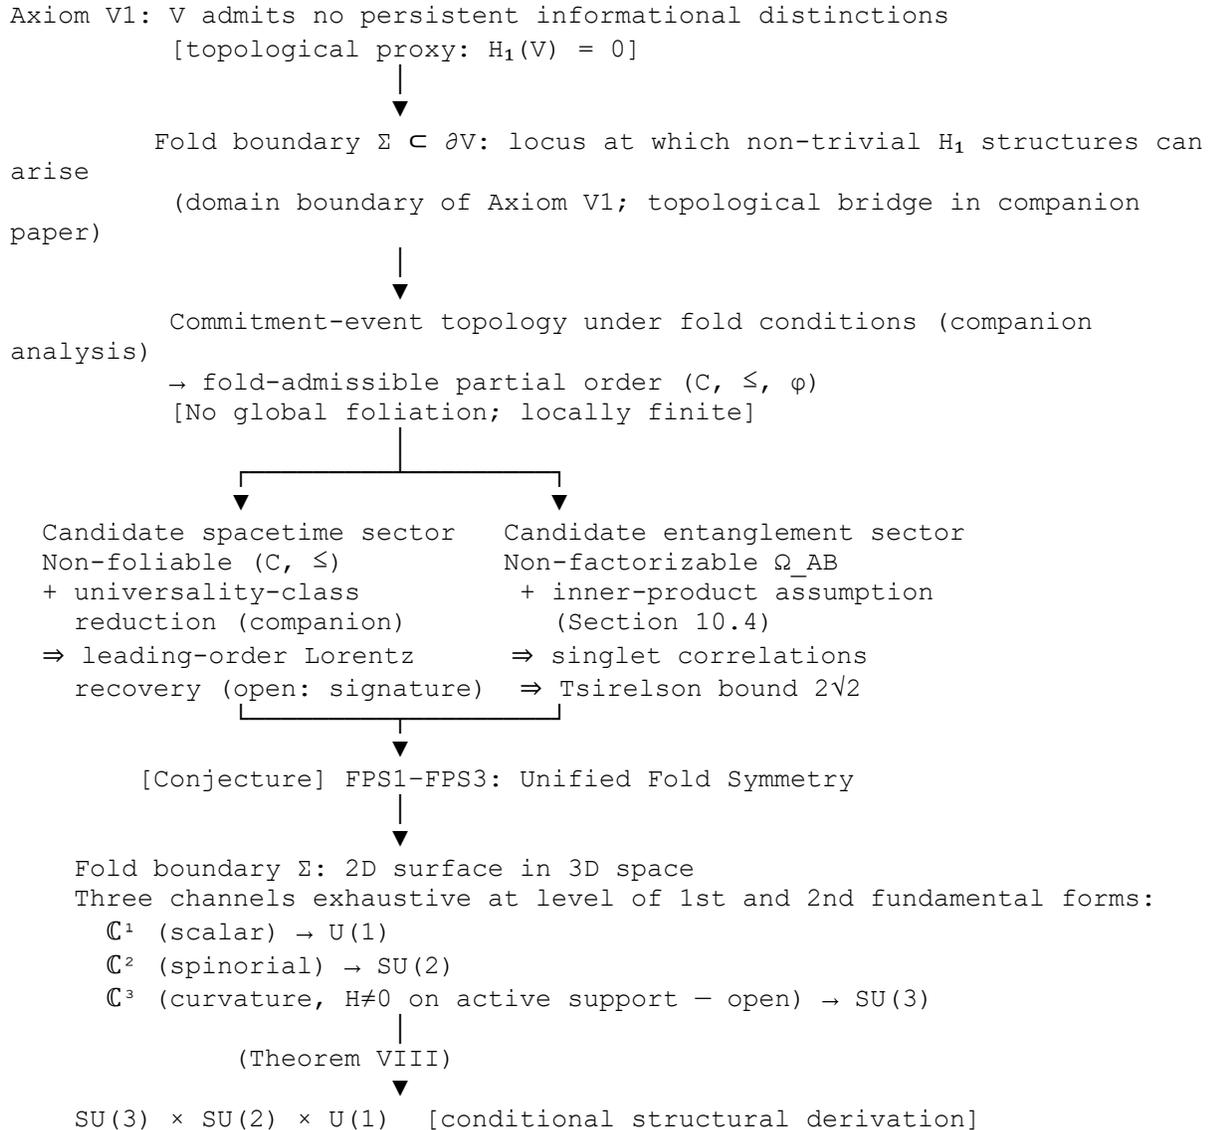
Discrete spacetime frameworks face a persistent challenge from Lorentz invariance. If spacetime has a discrete microscopic structure — a lattice, a network of nodes, a sequence of global time steps — then different inertial observers should decompose it differently. At sufficiently high boosts, the discreteness should become observable as a preferred frame: anisotropy in the speed of light, energy-dependent photon velocities, Lorentz-violating dispersion relations. Observations from the Fermi gamma-ray telescope have constrained such effects to extraordinary precision and found none.

A second challenge comes from quantum non-locality. Bell's theorem establishes that no theory assigning locally pre-existing values to measurement outcomes can reproduce quantum correlations. Any discrete commitment-based framework that simply pre-assigns outcomes to particles would be a local hidden variable theory — and Bell's theorem would immediately

falsify it. Quantum mechanics saturates the CHSH inequality at $2\sqrt{2}$ (the Tsirelson bound), and any adequate framework must reproduce this, neither falling short nor exceeding it.

Both objections are powerful — but both, as will be shown, rest on assumptions that VERSF does not share. The derivations presented here are structural rather than fully dynamical: they identify the minimal geometric and informational conditions required for these physical features to arise.

The framework can be summarised schematically as follows:



Sections 2–4 establish the pre-geometric structure (void, fold, topology of facts). Sections 5–8 derive the causal partial-order structure and its Lorentz consequences. Sections 9–11 address quantum correlations and their unification under FPS. Section 12 derives the gauge group structure from boundary geometry.

2. The Void

2.1 Definition

Axiom V1 — Zero-Distinction Substrate Let V denote the void substrate. V admits no persistent informational distinctions: any configuration of the substrate can be continuously transformed into any other without crossing an irreversible boundary. Formally, V contains no stable informational partitions. Equivalently, every configuration lies in a single reversibility-preserving connectivity class.

A useful topological signature of this condition is that the first homology group of V is trivial, $H_1(V) = 0$, which ensures that no non-contractible loops exist within V . This condition is weaker than the full informational statement of Axiom V1 but captures the absence of closed topological separations within the void.

Zero entropy here means not merely that the void is highly ordered, but that it is incapable of sustaining asymmetry of any kind. No distinguishable structure can be represented within it. The void is characterised not by content but by constraint.

2.2 What the Void Is Not

This does not require explaining why the void exists. Every foundational framework posits a substrate. The void is VERSF's substrate, characterised not by what it does but by what it cannot represent: no permanent distinction, no irreversible difference, no fact.

The void does not act. It does not shed asymmetry, restore symmetry, or drive dynamics. Rather: asymmetry cannot be represented in V . Distinguishable structure cannot persist within V . What appears in the framework as a fold is the structural consequence of this constraint — not something the void produces.

2.3 Reversibility as a Consequence

Because V admits no persistent asymmetry, all processes within it are reversible. This is not a dynamical law imposed on the void from outside — it is constitutive of Axiom V1. Reversibility is not a state V achieves; it is a property of what it means to be incapable of sustaining asymmetry.

This preempts the standard bootstrapping objection: there is no prior moment at which the void begins satisfying V1, no initiation of dynamics required. The void's symmetry is a boundary condition, not a process.

2.4 Topological Characterisation

As a topological regularity condition compatible with Axiom V1, we assume V is connected and has trivial first homology $H_1(V) = 0$. This does not define zero-distinction by itself; it merely excludes closed topological separations within the void. The full informational content of V is given by Axiom V1 in Section 2.1; the homology condition is a weaker proxy that captures one structural consequence:

$$H_1(V) = 0$$

This is the topological expression of V1. Zero entropy in topological terms means no non-trivial loop classes: no permanent separations, no facts, no closed regions from which there is no return. For simplicity we also assume V is simply connected ($\pi_1(V) = 0$), which guarantees that all non-trivial topology is externalised to the fold boundary Σ .

3. The Fold

3.1 Distinguishable Structure Cannot Persist in V

If V admits no distinguishable states, a question immediately arises: what becomes of distinguishable structure if it arises?

The answer is not annihilation but domain restriction. Because Axiom V1 defines the maximal domain in which distinguishable states cannot exist, any configuration containing such structure must lie outside the domain in which V1 holds. The fold boundary Σ therefore marks the minimal interface separating regions where V1 is satisfied from regions where distinguishable structure exists. This is a domain restriction argument, not a dynamical claim: Σ is where Axiom V1 ceases to apply, not where the void does something to eliminate asymmetry.

Asymmetry corresponds to distinguishable informational structure. Because such structure cannot be represented within V , it must appear at the boundary of V 's domain of validity. That boundary is the fold.

Lemma — Domain Boundary of Axiom V1 Let V denote the maximal domain on which Axiom V1 holds. Any configuration containing a persistent informational distinction cannot lie within V . Therefore such configurations lie in the complement of V . The interface separating these domains is the boundary: $\Sigma \subset \partial V$. This boundary is the fold. The fold therefore arises as the boundary of the maximal region in which the zero-distinction condition holds. It is not introduced dynamically but follows from the definition of the void as the largest domain satisfying Axiom V1.

Note on the topological bridge. This Lemma characterises Σ as a domain boundary but does not itself prove that Σ has non-trivial H_1 . A space can have trivial H_1 in the interior and trivial H_1 on its boundary simultaneously — contractible spaces with contractible boundaries are common.

The bridge from $H_1(V) = 0$ to $H_1(\Sigma) \neq 0$ requires the additional structure of commitment events (Section 4.2) and is established in detail in the companion paper *The Topological Threshold for Fact Formation*. The role of the present Lemma is to characterise why Σ is the correct locus for non-trivial topology to appear; the proof that non-trivial topology does appear there is deferred to that companion work.

Definition — Fold Boundary Σ The fold boundary Σ is the minimal interface separating V from regions containing irreversible informational structure. Formally, $\Sigma \subset \partial V$ satisfies: (i) on the V -side of Σ , $H_1 = 0$ (no persistent distinctions); (ii) on the non- V side of Σ , $H_1 \neq 0$ (permanent distinctions can exist); (iii) Σ has the minimum dimensionality consistent with (i) and (ii).

3.2 The Fold as Structural Consequence

Σ is not an event in time. It is the structural consequence of Axiom V1: a substrate with a maximal domain of validity excluding distinguishable structure must be bounded by an interface separating that domain from regions where such structure can exist. The fold does not arise from dynamics; it is what the existence of V and non- V together implies.

The fold has three properties:

- **Boundary, not region:** Σ separates void (topologically open, fully reversible) from non-void (topologically closed, capable of hosting irreversibility).
- **Minimal:** Σ represents the least structure sufficient to produce a permanent distinction.
- **Constitutive:** everything outside V is constituted by what accumulates at Σ . Physical structure is distinguishable content that cannot be represented within V and therefore exists only at or beyond Σ .

The discard region **D** is the region, adjacent to Σ on the non- V side, containing the alternatives that have been permanently committed away from. D satisfies the isolation and invariance conditions specified in Section 4.2. This boundary-localisation of physical structure is consistent with the general principle that realisable information is limited by boundary geometry rather than bulk degrees of freedom, a result that appears independently in the TPB area-scaling analysis of the VERSF program.

3.3 The Structural Inversion

This yields a significant reframing. The standard question — *how does order emerge from chaos?* — presupposes a positive origin story. VERSF's question is different: *given a substrate V that cannot represent asymmetry, what structures are necessarily present at its boundary?*

Under Axiom V1, Σ is structurally required, D is structurally required, and physical structure is the distinguishable content that accumulates at Σ — content that cannot be represented within V itself.

4. The Topology of Facts

4.1 What a Fact Requires

A fact is an irreversible commitment — a distinction that cannot be undone. For a fact to exist, there must be a structure in which two alternatives, once separated, cannot be recombined without loss.

This is a topological condition. Specifically, it requires a non-trivial first homology class (equivalently, a non-contractible loop class in π_1 whose abelianisation is non-trivial): two paths with the same endpoints that cannot be continuously deformed into one another. When such a separation exists, one path class can be permanently discarded. That discarding is the formation of a bit. That bit is a fact.

This condition is not merely conceptual but topological. As shown in the companion paper *The Topological Threshold for Fact Formation*, irreversible facts require the presence of a cyclic structure with non-trivial first Betti number $\beta_1 \geq 1$. Tree-like or one-dimensional structures cannot support local separability of alternatives: every branch can be recombined by backtracking. Only when a closed loop structure appears can two alternatives become permanently separated. Fact formation therefore coincides with the appearance of non-trivial topology in the substrate.

4.2 Why the Fold Creates Facts

The fold produces topological closure. At Σ , distinguishable structure accumulates that cannot be represented within V . This creates a discarded region D adjacent to Σ on the non- V side, satisfying two conditions:

- **Invariance:** the fold dynamics map D to itself — $T(D) = D$
- **Isolation:** no trajectory from D rejoins the active region under any allowed local operation

These two conditions together constitute the first non-trivial H_1 class. The boundary between D and the active region contains a non-contractible loop class — a closed path in the fold boundary topology that cannot be continuously contracted to a point. This is not merely a closed loop: contractibility must fail, which is guaranteed precisely by the isolation condition (no trajectory rejoins the active region). The fold boundary therefore contains a non-trivial homology class in H_1 .

This is the minimum topological structure for a fact to exist.

Whether these conditions are derived consequences of the fold dynamics or admissibility constraints that define what counts as a valid fold remains a question that the individual One Fold paper addresses in detail. Here they are treated as the operative definition of a commitment event.

4.3 Dimensional Signature

The topological threshold has a dimensional signature. The void (zero dimensions of factual structure) and a one-dimensional substrate both fail to support non-trivial H_1 : a line has no interior, no closed loops, no capacity for permanent separation. The minimum dimensionality that supports topological closure — and therefore facts — is two dimensions.

When multiple closure regions coexist, their mutual relationships — which regions are nested, adjacent, or causally separated — require an additional degree of freedom. Multiple closure regions cannot generically be embedded without intersection within a purely two-dimensional manifold; representing their adjacency and nesting relations generically requires embedding in a higher-dimensional ambient space. This embedding of closure relationships into a consistent geometric structure corresponds to a third dimension. The third dimension is therefore not a container for physics posited from outside: it is the geometric consequence of representing the relationships between coexisting factual closures. Two dimensions support minimal closure; three dimensions support the relational structure of multiple closures. A full proof that closure-region adjacency relations generically require a three-dimensional embedding remains to be given; a rigorous proof would require characterising the class of adjacency graphs generated by closure regions and showing that generic members of this class are non-planar. In the present work this dimensionality is treated as a structural consequence of the coexistence of multiple independent closure regions.

5. From Facts to Causal Structure

5.1 Commitment Events

Each topological closure — each fact — is a commitment event. The formation of a non-trivial H_1 class at the fold boundary is the moment at which one alternative is permanently discarded and the other is irreversibly registered.

A commitment event carries:

- A bit assignment $\varphi \in \{0,1\}$ — the fact fixed at this event
- A causal dependency relation — which prior commitment events were necessary for this one to occur

The set of all commitment events, ordered by causal dependency, constitutes the VERSF commitment structure.

5.2 The Commitment Structure

Proposition — Commitment Structure Given fold-admissible commitment events, the VERSF commitment structure is a triple (C, \leq, φ) where:

- C is a countable set of commitment events
- \leq is a partial order on C encoding causal precedence. The reflexive closure of the causal relation is used throughout; the strict causal relation is obtained by removing reflexivity ($c < d$ iff $c \leq d$ and $c \neq d$).
- $\varphi: C \rightarrow \{0,1\}$ is the bit assignment

The partial order satisfies the standard conditions:

- **Reflexivity:** $c \leq c$ for all $c \in C$
- **Antisymmetry:** $c \leq d$ and $d \leq c$ implies $c = d$
- **Transitivity:** $c \leq d$ and $d \leq e$ implies $c \leq e$

Crucially, the partial order does not require every pair of events to be comparable. Two events c and d may satisfy neither $c \leq d$ nor $d \leq c$. Such events are spacelike in the emergent sense: causally unrelated, with no fact-formation precedence between them.

5.3 Fold Admissibility

Not every partial order on commitment events is consistent with the void's symmetry constraint. The VERSF commitment structure is restricted to *fold-admissible* partial orders. C is countably infinite — discrete without being finite. The admissibility conditions are:

Local finiteness: for any two events $c, d \in C$, the interval

$$[c, d] = \{e \in C : c \leq e \leq d\}$$

is finite. This encodes discreteness without implying foliation.

No global foliation: the commitment structure admits no global foliation compatible with the partial order. That is, there exists no partition of C into equivalence classes $\{S_t\}$ of mutually incomparable events such that the S_t are totally ordered by \leq between classes and every event lies in exactly one class. The commitment structure does not admit a decomposition into spatial hypersurfaces consistent with the causal order.

Fold consistency: every commitment event corresponds to a non-trivial H_1 class in the fold boundary topology, and the precedence relation \leq is induced by topological containment of discard regions — specifically, $c \leq d$ if and only if the discard region D_c is topologically contained in the causal past of D_d .

The no-global-foliation condition is both required for the Lorentz argument and expected to follow from the fold topology, since the void has no preferred temporal direction and its symmetry constraint does not single out any spatial hypersurface. Whether it is fully derived from fold consistency or requires independent statement is a question that needs explicit resolution in the One Fold paper.

5.4 The Partial Order Is Not a Foliation

The key property of (C, \leq, φ) for what follows is that it carries causal order without carrying foliation. A partial order does not define a stack of universal moments. There is no "time slice" of the commitment structure, no global "now," no partition of C into spatial hypersurfaces. Causal precedence between individual events does not generate simultaneity between others.

This is what distinguishes VERSF's discrete structure from a discretised spacetime manifold and, consequently, what dissolves the standard Lorentz objection.

6. The Lorentz Objection and Its Hidden Assumption

6.1 The Standard Argument

The standard objection to discrete spacetime runs as follows: if spacetime has a discrete microscopic structure, different inertial observers moving at different velocities relative to that structure will decompose it differently. At sufficiently high boosts, the discreteness should become observable as a preferred frame — anisotropy in the speed of light, Lorentz-violating dispersion relations. Fermi telescope observations have constrained such effects to extraordinary precision and found none.

6.2 The Hidden Assumption

The objection is powerful, but it is only powerful against frameworks that share a specific hidden assumption:

There exists a microscopic global simultaneity structure that different observers cut differently.

The objection imagines boosted observers slicing an underlying crystal at different angles. It presupposes that the substrate has the structure of a discretised spacetime manifold: nodes arranged in a background container, spatial layers stacked in time. The discreteness is *spatial* — nodes in a three-dimensional lattice — and the Lorentz problem arises because the lattice spacing defines a preferred frame.

VERSF does not share this assumption, for two independent reasons:

The substrate is non-foliable. The void has no global simultaneity structure. The commitment structure (C, \leq, φ) explicitly excludes any partition of C into spatial hypersurfaces compatible with \leq . There is no universal "now" to slice.

Discreteness is not spatial. What is discrete in VERSF is the formation of irreversible facts — commitment events. A countable set of causally ordered events does not automatically define a preferred spatial slicing. Causal order can be discrete without becoming a stack of universal moments.

The standard objection therefore does not apply to VERSF. The objection needs a lattice to break. There is no lattice.

7. How Lorentz Invariance Emerges

Given the commitment structure (C, \leq, φ) , the task is to show that its invariant relational statistics yield Minkowski interval behaviour under coarse-graining for inertial observers. This section identifies the key steps.

7.1 The Interval from Commitment Counting

In causal set theory, the Lorentzian spacetime interval between two causally related events is proportional to the number of causal set elements between them (the Myrheim-Meyer result). For events $c \leq d$:

$$|\{e \in C : c \leq e \leq d\}| \sim \text{Vol}(\text{bigl}(J^+(c) \cap J^-(d)\text{bigr})$$

where the volume is measured in the emergent Minkowski spacetime.

The target for VERSF is an analogous relation: the Lorentzian interval between two commitment events should be encoded in the count of intermediate commitments. Because this count is an invariant of the partial order — not dependent on any observer's slicing — it is automatically Lorentz invariant if the relation holds. No observer can cut the substrate differently because there is no spatial layer structure to cut. Similar interval-counting relations appear in causal set theory, where Lorentz invariance emerges statistically from the absence of preferred embeddings (Bombelli et al. 1987; Henson 2006): the Poisson sprinkling distribution is itself Lorentz invariant, and no preferred frame can be read off from a random causal set.

In causal set theory, the Myrheim-Meyer result depends on Poisson sprinkling — a random embedding of causal set elements into Minkowski space with Poisson-distributed density. In VERSF, fold admissibility plays the structural role that Poisson sprinkling plays in causal set theory: it is the constraint that determines which partial orders are realised and, consequently, which statistical properties their intervals inherit. Whether fold admissibility enforces the interval-counting relation is precisely the open question. The Myrheim-Meyer interval-counting relation is therefore the target behaviour for VERSF, not a justification of the approach.

7.2 Light Cones from Upper-Set Boundaries

A light cone in Minkowski spacetime is the boundary of the causal future of an event. In the commitment structure, the causal future of c is:

$$\{\uparrow c\} = \{d \in C : c \leq d\}$$

The conjecture is that the boundary $\partial(\uparrow c)$, in the appropriate coarse-graining limit, has the geometry of a Minkowski light cone. This requires showing that fold-admissible partial orders have the right density and connectivity properties for this boundary to be cone-shaped — neither a sphere nor a fractal — which is a non-trivial constraint on the class of admissible partial orders.

7.3 Time Dilation from Chain Length

Relativistic time dilation must emerge from the structure of paths in the partial order. A maximal chain in the partial order — a sequence $c_1 \leq c_2 \leq \dots \leq c_n$ with no intermediate events — corresponds to a timelike worldline. The proper time along that worldline is proportional to chain length $|\text{chain}|$.

For time dilation to emerge correctly, paths that traverse greater spatial separation through the commitment structure must correspond to shorter chains. Formally, for two maximal chains γ_1, γ_2 connecting c to d with γ_2 having greater spatial extent:

$$|\gamma_1| > |\gamma_2|$$

This is not self-evident and constitutes a non-trivial constraint on fold-admissible partial orders. A plausible mechanism is that fold admissibility enforces a density gradient — commitment events are sparser along spatially extended paths — but this needs explicit derivation.

7.4 No Observable Preferred Frame

The strongest constraint is the absence of any observable preferred frame. In the commitment structure, this requires showing that no invariant of (C, \leq, φ) defines a preferred decomposition into spatial layers.

The partial order admits no global foliation — no decomposition into spatial hypersurfaces compatible with \leq — so there is no sense in which any observer's "now" picks out a distinguished subset of C . All inertial observers see the same partial order, coarse-grain it in different ways, and, if the framework is correct, extract the same Lorentzian interval statistics. The frame-dependence is entirely in the coarse-graining, not in the underlying structure.

7.5 Leading-Order Lorentz Recovery from Fold-Compatible Dynamics

The preceding sections frame Lorentz emergence as an open problem — whether fold admissibility selects partial orders with Lorentzian interval statistics. The companion analysis in *The One Fold* provides a concrete mechanism at the level of effective dynamics that turns the open problem into a sharply stated reduction target.

Companion Reduction Result — Leading-Order Lorentz Recovery from Fold-Compatible Dynamics Consider a local, Hermitian, translation-invariant nearest-neighbour Hamiltonian on a cubic lattice with four-component internal fiber (the minimal fold quantum realization of Section 10.4). If its momentum-space kernel admits the low-momentum expansion: $K(\mathbf{p}) = v \cdot \boldsymbol{\alpha} \cdot \mathbf{p} + O(a^2 |\mathbf{p}|^3)$ where v is an effective speed

and α are Dirac-like matrices, then the low-energy dispersion satisfies: $E^2 = v^2 |\mathbf{p}|^2 + O(a^2 |\mathbf{p}|^4)$ Consequently, the effective long-wavelength sector is Lorentz invariant at leading order, with violations suppressed by $O((E/E^*)^2)$ where E^* is the fold discreteness scale. The detailed construction and verification of this reduction are given in the companion One Fold analysis; the role of the present section is to state the universality-class consequence for the VERSF program. This is a result within the companion fold-realization model, not yet a theorem about all fold-admissible commitment structures.

This result establishes that fold-compatible dynamics with a four-component internal fiber fall into a universality class that recovers Lorentz invariance at leading order. The remaining question is whether fold-admissible commitment structures necessarily flow into this universality class under coarse-graining — transforming the problem from "can any fold-compatible discrete structure produce Lorentz covariance?" (answered yes) to "does fold admissibility select this universality class?" (the precise remaining open problem).

8. Relation to Causal Set Theory

The VERSF commitment structure closely resembles causal set theory (Bombelli et al. 1987; Sorkin 1991, 2003) in formal outline but differs in two important respects.

Origin. Causal sets posit the partial order as a primitive, then seek dynamical laws (such as the classical sequential growth model of Rideout and Sorkin 2000) that generate the right statistical properties. VERSF derives the partial order from the fold topology, which is itself a consequence of the void's symmetry constraint (Axiom V1). The partial order is not a starting point but a derived object.

Constraint. Causal set theory relies on the Hauptvermutung conjecture — that generic random partial orders sprinkled into Minkowski space yield Lorentzian geometry at large scales. VERSF has the additional resource of fold admissibility, which restricts the partial order to those consistent with the void topology. The question is whether this constraint is strong enough to select Lorentzian statistics without additional assumptions, and whether the admissible class is neither so restrictive that it excludes all Lorentzian geometries nor so permissive that it collapses back to a genericity argument. The present work does not claim to solve the causal set Hauptvermutung problem; rather it proposes that fold admissibility provides an additional structural constraint that may narrow the class of admissible partial orders.

This distinction matters. The Lorentz emergence problem in VERSF is a sharper problem than the one faced by causal set theory: *does fold admissibility alone, without Poisson sprinkling, select partial orders with the right interval statistics?* It is also potentially more tractable, because the admissibility constraint comes from a specific topological structure rather than being chosen for convenience.

Where the frameworks agree. Both frameworks treat the partial order as fundamental and spacetime as emergent. Both use the Myrheim-Meyer interval-counting relation as a target for

Lorentz covariance. Because Poisson sprinkling is itself Lorentz invariant, no preferred frame can be extracted from a random causal set — a result that situates the no-preferred-frame argument historically and shows it is not specific to VERSF. The VERSF contribution is to provide a derivation of why the partial order has the structure it does, rather than positing it.

9. Non-Factorizable Commitment and Bell Non-Locality

9.1 Why Local Bit Assignment Fails

Bell's theorem establishes that no theory assigning locally pre-existing values to measurement outcomes can reproduce quantum correlations. Formally, for any local hidden variable model in which each particle in an entangled pair carries a hidden variable λ drawn from a distribution $\rho(\lambda)$ independent of measurement settings, with outcomes:

$$A = A(a, \lambda), B = B(b, \lambda)$$

the CHSH combination of correlations satisfies:

$$|S| \leq 2$$

No choice of λ , $\rho(\lambda)$, or response functions can exceed this bound while maintaining locality and setting-independence. Quantum mechanics reaches $|S| \leq 2\sqrt{2}$.

If VERSF's commitment structure simply pre-assigned outcome bits to each wing of an entangled pair — treating entanglement as a correlation between two independently committed facts — it would be a local hidden variable theory. Bell's theorem would apply and the framework would be falsified by experiment. The resolution requires identifying what the void substrate assigns instead of local pre-existing bits.

9.2 The VERSF Alternative: Non-Factorizable Commitment

Assumption — Non-Factorizable Joint Commitment For an entangled pair, the void substrate does not assign independent hidden states to each wing. Instead, it assigns a single non-factorizable substrate relation Ω_{AB} , satisfying $P(A, B | a, b, \Omega_{AB}) \neq \int d\lambda \rho(\lambda) P(A|a, \lambda) P(B|b, \lambda)$.

This assumption is the VERSF alternative to local hidden variables. Its motivation is that the void substrate has no preferred frame (Section 5.4), no global slicing, and no mechanism for assigning spatially separated pre-existing values that are consistent with Lorentz covariance. Whether this assumption can be derived from fold admissibility alone — rather than posited — is the entanglement sector of the open derivation program in Section 15.

Before commitment, the pair does not consist of two separate local value carriers. It exists as a single distributed potential relation across the void substrate. Spacelike separation is a property

of the emergent spacetime description, not of the substrate itself. When measurements occur at sites A and B, those two local commitment events are separate only in emergent spacetime. At the substrate level, they are two local manifestations of one shared pre-geometric constraint.

This is directly continuous with the partial order picture of Section 5. Two spacelike-separated commitment events $c, d \in C$ — satisfying neither $c \leq d$ nor $d \leq c$ — are causally unrelated in the emergent sense, but may share a common substrate relation Ω not captured by the causal order alone. Entanglement is the persistence of such a shared relation into the spacelike-separated commitment regime.

9.3 Settings as Commitment Directions

Local measurement settings \mathbf{a} and \mathbf{b} correspond to unit vectors in an emergent measurement space $\mathbf{a}, \mathbf{b} \in S^2$. The entangled pair corresponds to a substrate state with a singlet-like anti-alignment constraint. Measurement operators are:

$$\hat{A}(\mathbf{a}) = \mathbf{a} \cdot \boldsymbol{\sigma}, \hat{B}(\mathbf{b}) = \mathbf{b} \cdot \boldsymbol{\sigma}$$

where $\boldsymbol{\sigma}$ are Pauli-like generators of the commitment algebra. The correlation function is:

$$E(\mathbf{a}, \mathbf{b}) = \langle \omega | \hat{A}(\mathbf{a}) \otimes \hat{B}(\mathbf{b}) | \omega \rangle = -\mathbf{a} \cdot \mathbf{b}$$

This is the quantum singlet correlation; this form is derived from structural constraints on the void substrate in Section 10. What a measurement setting *is* in VERSF terms: a local orientation of commitment basis — the direction along which the apparatus resolves the shared substrate relation into a definite fact. What measurement *does*: it resolves the shared relational mode into paired facts according to a compatibility law. Outcomes are locally random, jointly constrained, non-factorizable, and non-signalling.

10. The Tsirelson Bound from Commitment Geometry

10.1 Five Constraints on the Void Substrate

A more VERSF-native derivation proceeds from five constraints on how the void substrate treats entangled pairs, each motivated directly by the void's defining properties.

Constraint 1 — Unbiased local marginals: $P(A = \pm 1 | \mathbf{a}) = 1/2$ and $P(B = \pm 1 | \mathbf{b}) = 1/2$. No local measurement is biased. The void substrate admits no persistent asymmetry, so it cannot prefer one outcome at either wing.

Constraint 2 — Relational setting dependence: $P(A, B | \mathbf{a}, \mathbf{b}) = F(AB, \mathbf{a} \cdot \mathbf{b})$. Joint probabilities depend only on the relative geometry between local measurement contexts, not on their absolute orientations. This follows from the rotational invariance of the void.

Constraint 3 — Rotational invariance: Only the angle between settings matters. The void substrate has no preferred direction, so the compatibility law cannot privilege any absolute orientation.

Constraint 4 — Anti-correlation at equal settings: $E(a, a) = -1$. When both wings measure in the same direction, outcomes are perfectly anti-correlated. This is the singlet constraint encoding the anti-alignment of the joint substrate relation.

Constraint 5 — No signalling: Marginals at A do not depend on b, and vice versa. Local commitment events cannot transmit information, consistent with the non-foliable structure of (C, \leq, φ) .

10.2 Deriving the Singlet Correlation Law

Among functions $F(AB, \mathbf{a} \cdot \mathbf{b})$ satisfying constraints (1–5), the minimal solution is defined as the one with the lowest polynomial degree in the invariant $\mathbf{a} \cdot \mathbf{b}$. Restricting to functions analytic in the invariant $\mathbf{a} \cdot \mathbf{b}$ — as required by the rotational invariance of Constraint 3, which forces dependence only on this scalar invariant — the lowest-order admissible term is linear.

Analyticity is assumed here as the weakest regularity condition compatible with smooth variation of measurement context. Higher-order terms introduce additional tunable parameters not required by the constraints and therefore represent additional relational structure beyond what the constraints demand. Such parameters would represent additional observable relational structure beyond the five constraints, contradicting the minimality criterion. The lowest-order admissible solution is linear in $\mathbf{a} \cdot \mathbf{b}$.

The minimal functional form consistent with all five constraints is:

$$P(A, B | \mathbf{a}, \mathbf{b}) = \frac{1}{4}(1 - AB \cdot \mathbf{a} \cdot \mathbf{b})$$

The correlation function is then:

$$E(a, b) = \sum_{\{A,B\}} AB \cdot \frac{1}{4}(1 - AB \cdot \mathbf{a} \cdot \mathbf{b}) = -\mathbf{a} \cdot \mathbf{b}$$

This is the quantum singlet correlation, derived from VERSF structural constraints together with the inner-product assumption of Section 10.4.

10.3 The Tsirelson Bound

The CHSH combination is:

$$S = E(a, b) + E(a, b') + E(a', b) - E(a', b')$$

Using $E(x, y) = -\mathbf{x} \cdot \mathbf{y}$, the maximum is achieved at the standard 45° geometry between settings, giving:

$$|S| \leq 2\sqrt{2}$$

This is the Tsirelson bound (Tsirelson 1980; Cirel'son 1980). The framework reaches quantum correlations exactly.

The derivation here proceeds from the five structural constraints of Section 10.1 rather than from operator algebras. Tsirelson's original proof establishes the bound from the representation theory of operator algebras acting on Hilbert space (Cirel'son 1980); his subsequent work (Tsirelson 1987) and Landau (1987) give tighter formulations using operator norms and correlation matrices. Here the bound emerges from the anti-alignment constraint plus the inner-product geometry of FPS3. The two routes are complementary: the operator-algebraic proof is rigorous; the VERSF derivation is motivated but requires the inner-product assumption stated in Section 10.4.

10.4 Why Superquantum Correlations Are Excluded

Proposition — Minimal Quantum Realization of Commitment States If a commitment event carries one binary distinction together with reversible Z_2 directionality, then its minimal reversible quantum realization is a complex Hilbert space $\mathbf{H_fold} \cong \mathbb{C}^4$. Physical states are rays in $\mathcal{P}(\mathbf{H_fold}) = \mathbb{C}P^3$, and the natural composition law for local amplitudes is therefore inner-product based. This is the structure established in the companion paper *The One Fold* and adopted here as the quantum-substrate premise for the VERSF framework.

Assumption — Complex Inner-Product Structure (adopted from minimal quantum realization) Commitment amplitudes inhabit a complex inner-product space, inherited from the $\mathbf{H_fold} \cong \mathbb{C}^4$ realization above. The complex structure is required to represent phase relations in commitment amplitudes and to permit non-trivial unitary transformations between measurement contexts. This assumption introduces Hilbert-like structure into the commitment algebra and gives local commitment operators bounded norm $\|\hat{A}\| \leq 1$.

Honest limit. This does not derive the complex inner product from the void substrate alone. It derives it from the void-fold framework together with the quantum-substrate premise that commitment states admit minimal reversible quantum realization. The inner-product structure is therefore inherited from the companion fold analysis, not imported ad hoc at the Bell stage.

A structural motivation for why complex rather than real amplitudes arise connects to information geometry. Complex structure is the minimal algebra supporting reversible composition and phase relations under binary commitment: real amplitudes permit reversible combinations but cannot represent the full interference structure that emerges when a commitment state is rotated relative to a measurement context. Wootters (1981) and Braunstein and Caves (1994) establish that complex Hilbert space is distinguished among inner-product spaces by its statistical distance structure — the Fubini-Study metric on $\mathbb{C}P^n$ is the unique metric arising from a statistically natural measure on quantum states. This suggests that the complex structure, rather than being an arbitrary input, is selected by the requirement that commitment amplitudes define a statistically well-behaved geometry on the space of possible commitment outcomes. Deriving this selection directly from Axiom V1 remains the deepest open problem of the framework.

The inner-product structure constrains the CHSH operator norm to exactly $2\sqrt{2}$. FPS3 (Section 11.2) is the projection geometry postulate that enforces this norm bound operationally. The complex inner-product structure is therefore foundational to both the Tsirelson derivation and the gauge-channel complexification of Section 12. Its origin within the void–fold axioms remains an open problem, stated at the top of Section 15. Whether the complex inner-product structure can be derived from the void substrate alone — without the quantum-substrate premise — is the deepest remaining foundational gap.

10.5 The Residual Gap

The derivation above is a proposed mechanism, not yet a full derivation from deeper VERSF axioms. The honest accounting is as follows.

Established: Classical local bit commitment cannot reproduce quantum correlations. VERSF requires a non-factorizable joint commitment amplitude. Five natural constraints on the void substrate's treatment of entangled pairs select the singlet correlation law as the minimal compatible solution. The singlet correlation law immediately yields the Tsirelson bound.

Not yet established: Why the void substrate naturally generates inner-product structure for commitment amplitudes. Whether the five constraints follow from the fold topology and void symmetry constraint, or must be posited independently. Whether the commitment algebra's inner-product geometry is the same structure that generates Lorentzian geometry of the emergent partial order.

Section 11 addresses this residual gap directly. The Fold Projection Symmetry Principle provides a single framework from which both the Lorentz and Tsirelson results follow, and identifies precisely what must be derived to close the gap between mechanism and theorem.

11. Fold Projection Symmetry: The Common Root

11.1 The Shared Structure of Both Problems

Sections 7 and 10 addressed the Lorentz and Tsirelson problems by separate routes. This section shows they have the same deep shape and identifies the single principle responsible for both.

In each case, the challenge is the same: the fold must permit relational structure to become observable while forbidding observers from accessing hidden substructure. Observers reach the fold only through local projections — measurement orientations, inertial frames — and what they can extract is constrained. The two constraints are:

- No observer can extract a preferred inertial slicing from the fold. Lorentz covariance follows.
- No pair of observers can extract more relational information than the fold's norm geometry permits. Tsirelson follows.

These are not two separate restrictions. They are the same restriction applied in two observational contexts. The fold permits only invariant quadratic relational projections, and conceals everything else. This is the Fold Projection Symmetry Principle.

11.2 The Fold Projection Symmetry Principle

Let the fold carry a pre-geometric relational object F . Let local measurements and inertial descriptions correspond to choosing local projection contexts u, v, \dots on F . The governing principle is:

Fold Projection Symmetry Principle (FPS): All observable commitment statistics depend only on invariant relational overlaps between local commitment contexts, never on any underlying foliation, embedding, or hidden coordinate description of the fold.

This is formalised through three axioms.

Axiom FPS1 — Non-sliceability. There exists no globally observable foliation of F . The fold admits no decomposition into simultaneous spatial layers that any observer can detect. This is exactly the non-foliability of the commitment structure established in Section 5.4.

Axiom FPS2 — Contextual projection covariance. Observable outcomes depend only on invariant overlap data between projection contexts, not on any absolute labelling of the fold. If two local descriptions are related by a symmetry of F , they produce identical observable statistics.

Postulate FPS3 — Minimal Relational Composition. Let observable commitment statistics depend only on overlaps between local projection contexts, as required by FPS2. Among admissible composition laws for these overlaps, the physically admissible law is the weakest one that: (i) preserves bounded binary outcomes; (ii) introduces no additional invariantly measurable label beyond the local contexts themselves; and (iii) is stable under composition of local projections. The minimal such law is quadratic:

$$\langle \Gamma(u, v) \rangle = \langle \text{lange } u, v \rangle, \quad |\Gamma(u, v)| \leq 1$$

with local observables represented by bounded self-adjoint operators with spectrum $\{\pm 1\}$.

FPS3 is not yet derived from fold topology alone. Its present status is a constrained admissibility principle: among projection-composition laws compatible with FPS1–FPS2, the quadratic norm law is the minimal one that does not introduce additional invariantly measurable relational structure. Within the present framework, superquantum composition laws are treated as requiring additional accessible relational structure beyond the projection contexts themselves — and FPS2 already requires that no such additional invariant label appear. FPS3 selects the minimal composition law consistent with this requirement.

FPS1 and FPS2 together prevent preferred frames. FPS3 prevents superquantum correlations by selecting the minimal composition law that survives the no-extra-label requirement. The same

quadratic constraint appears in both sectors — this is the structural content of the Unified Fold Symmetry Conjecture.

Open problem. FPS3 is a constrained admissibility principle, not yet a theorem derived from void topology alone. A full derivation would establish that the void's symmetry constraint (Axiom V1) — which forbids extra distinguishability labels within V — propagates to the fold boundary to forbid superquantum composition laws at Σ . Until this derivation is complete, FPS3 is the best-motivated minimal postulate, not a consequence.

11.3 The Spacetime Sector: Lorentz Covariance

Suppose the coarse-grained separation between commitment events is encoded by a quadratic invariant arising from fold projection covariance:

$$I(x) = g_{\mu\nu} x^\mu x^\nu$$

By FPS2, this bilinear form is the same for all inertial observers — it is an invariant of the fold, not of any particular observer's projection context. The transformations that preserve $I(x)$ form the effective spacetime symmetry group. The route is:

1. FPS1: no observable slicing of the fold exists
2. FPS2: only invariant interval-like relations survive coarse-graining
3. The symmetry group preserving those relations is determined by the signature of the preserved quadratic form
4. For signature $(+, -, -, -)$, this group is the Lorentz group
5. Therefore no preferred frame appears in the effective description

The remaining open task — showing that the fold topology produces a coarse-grained bilinear form with Minkowski rather than Euclidean signature — is the same problem identified in Section 7, now embedded within the unified FPS framework.

11.4 The Entanglement Sector: The Tsirelson Bound

For entangled commitment events, local outcomes arise from projecting a shared fold relation F into local contexts a and b . By FPS2, the correlation function depends only on the invariant overlap:

$$E(a, b) = -\Gamma(a, b) = -\langle a, b \rangle = -\mathbf{a} \cdot \mathbf{b}$$

This is the quantum singlet correlation, recovered from FPS2 alone together with the anti-alignment of the shared fold relation.

Proposition: If $E(a, b) = -\langle a, b \rangle$ and the pairing $\langle \cdot, \cdot \rangle$ satisfies the quadratic norm law of FPS3, then the CHSH combination satisfies $|S| \leq 2\sqrt{2}$.

The local commitment operators are assumed to be bounded self-adjoint operators with spectrum $\{+1, -1\}$, consistent with binary commitment outcomes.

Proof sketch (assuming the standard bounded-operator representation of the local commitment algebra): Write the CHSH operator as:

$$\hat{B} = \hat{A}(a) \otimes \hat{B}(b) + \hat{A}(a) \otimes \hat{B}(b') + \hat{A}(a') \otimes \hat{B}(b) - \hat{A}(a') \otimes \hat{B}(b')$$

By FPS3, each local operator has eigenvalues ± 1 and satisfies $\|\hat{A}\| \leq 1$. Computing \hat{B}^2 :

$$\hat{B}^2 = 4I - [\hat{A}(a), \hat{A}(a')] \otimes [\hat{B}(b), \hat{B}(b')]$$

The commutator terms contribute at most 4 to the operator norm, giving $\|\hat{B}\|^2 \leq 8$, hence:

$$\|\hat{B}\| \leq 2\sqrt{2} \blacksquare$$

The Tsirelson bound follows directly from FPS3 — the same norm-bounded quadratic projection law that is also responsible for the quadratic invariant in the spacetime sector. The two sectors share not just an analogy but the same axiom.

11.5 The Unified Fold Symmetry Conjecture

The two sectors share a non-sliceable relational structure F , local projection contexts as the only observable interface, and a preserved quadratic form as the invariant of the fold symmetry group. This convergence motivates the following:

Conjecture — Unified Fold Symmetry Let the fold be a non-sliceable pre-geometric relational structure whose observable commitments are generated only through invariant norm-bounded projections (FPS1–FPS3). Then: (i) no observable preferred foliation exists, yielding effective Lorentz covariance in the coarse-grained event algebra; and (ii) no admissible joint commitment statistics exceed the norm bound induced by the same projection geometry, yielding the Tsirelson limit for Bell-type correlations. Both results are representations of a single fold symmetry group G_F acting on the relational space H_F of the fold — one in the spacetime sector, one in the entanglement sector — both preserving the same quadratic invariant.

This is a conjecture, not a theorem. The three-part derivation program needed to establish it is given in Section 11.6.

11.6 What a Full Proof Would Require

Three things are needed to close the conjecture.

A single fold symmetry object. Define the fold symmetry group G_F acting on the pre-geometric relational space H_F . This group should be derivable from the void symmetry constraint and the fold topology — not postulated independently. The void's defining property

(admitting no persistent asymmetry) constrains which transformations leave the fold structure invariant, and G_F is the group of those transformations.

Two sectors as representations of G_F . Show that one representation of G_F on H_F generates the coarse-grained spacetime interval structure with Minkowski signature; and that another representation generates the measurement-context overlap structure with the quadratic correlation kernel $E(\mathbf{a}, \mathbf{b}) = -\mathbf{a} \cdot \mathbf{b}$. These are to be derived, not assumed.

One quadratic invariant. Show there is a unique or natural quadratic invariant preserved by G_F , and that in the spacetime representation it becomes Minkowski interval structure, while in the entanglement representation it becomes the bounded correlation kernel producing the Tsirelson bound.

If all three are established, Lorentz invariance and the Tsirelson bound are not two separate features of the framework requiring separate justification. They are the same feature of the fold geometry, expressed in two different observational contexts.

12. The Standard Model Gauge Group from Fold Boundary Geometry

Scope and caveats. This section derives $SU(3) \times SU(2) \times U(1)$ as the gauge group structure consistent with fold boundary geometry. Three limitations should be noted at the outset: (1) the $H \neq 0$ condition is structurally motivated but not yet formally derived from $\mathcal{S}_{\text{fold}}$; (2) the derivation selects $SU(2)$ but does not establish its chirality — $SU(2)_L$ requires a separate argument; (3) gauge dynamics, coupling constants, and the Higgs mechanism are not addressed. The derivation establishes the *group structure*, not the full gauge theory.

12.1 The Prior State of the Derivation

Theorem VIII of the VERSF gauge group paper establishes the following:

Given complex vector spaces of dimensions 1, 2, and 3 carrying faithful irreducible unitary representations of minimal compact connected Lie groups, the groups are uniquely $U(1)$, $SU(2)$, and $SU(3)$, and their direct product is $SU(3) \times SU(2) \times U(1)$.

Irreducibility reflects the assumption that each geometric channel represents a single elementary degree of freedom rather than a composite of independent subsystems.

This mathematical result is rigorous. What remained conjectural was why exactly those three dimensions should be assigned to the geometric channel types at the fold boundary. The present section closes that gap. The (1, 2, 3) assignment is derived from three facts established elsewhere in this paper: the fold creates a 2-dimensional boundary (Section 4.3); commitment amplitudes

require a complex inner-product space (Section 10); and the fold boundary is an active commitment interface, not a passive equilibrium surface (derived below in Section 12.6).

12.2 The Three Geometric Channels

At any point p on the fold boundary, the local geometric information available to a commitment event decomposes into exactly three distinct types.

Theorem (Fundamental Theorem of Surface Theory — Gauss-Codazzi). A 2-dimensional smooth surface Σ embedded in 3-dimensional Riemannian space is locally determined, up to rigid motion, by its first fundamental form (metric induced on Σ) and its second fundamental form (shape operator encoding extrinsic curvature). These two forms are necessary and sufficient for local characterisation.

Consequently, local geometric information at a point on Σ decomposes exhaustively into: (i) scalar separation, encoded by the metric; (ii) local orientation of the surface normal, encoding directional structure; and (iii) extrinsic curvature, encoded by the shape operator. The step from two fundamental forms to three channels is as follows: the first fundamental form contributes the scalar separation channel (the metric determines distances); the second fundamental form contributes both the orientation channel (via the direction of the surface normal) and the curvature channel (via the independent components of the shape operator). No additional independent local invariant exists at the level of first and second fundamental forms.

The three channels therefore are not a motivated choice among alternatives. They are the exhaustive decomposition of local fold boundary geometry at the level of first and second fundamental forms, given by the fundamental theorem of surface theory. This geometric decomposition corresponds to the same three structural sectors identified independently in the VERSF structural analysis program, where the excitation space naturally splits into a scalar sector, a two-component orientation sector, and a three-dimensional interaction sector, reinforcing the 1–2–3 structure across the framework.

12.3 Channel 1: Scalar Separation $\rightarrow U(1)$

The first channel answers a purely scalar question: how many intermediate commitment events lie in the partial order interval between two nearby events? This encodes separation — proximity in the causal structure — and is invariant under all local rotations of the fold boundary. It requires no directional data and no shape data. A single complex number is both necessary and sufficient. The scalar separation channel becomes complex not because separations themselves are complex numbers, but because commitment amplitudes associated with separations inherit the complex structure of the commitment algebra established in Section 10.

The space is \mathbb{C}^1 . By Theorem VIII, the minimum compact connected Lie group acting faithfully and irreducibly on \mathbb{C}^1 is $U(1)$. ✓

This step requires only that the fold produces a boundary at all. The $U(1)$ factor is an immediate consequence of the void possessing any boundary whatsoever.

12.4 Channel 2: Spinorial Orientation \rightarrow SU(2)

The second channel concerns the local orientation of the fold boundary — specifically, the direction the fold normal points as it varies across the 2-dimensional surface embedded in the emerging 3-dimensional structure established in Section 4.3.

Orientation in 3D has spinorial character. A rotation of the normal vector by 2π returns it to its original direction, but the corresponding spinorial object acquires a phase of -1 . A single complex number cannot faithfully represent this: orientation information in 3D is irreducibly 2-component. The minimum complex vector space faithfully representing a spinor in 3D is \mathbb{C}^2 . This is a theorem of SU(2) representation theory: the fundamental representation of SU(2) is 2-dimensional, and it is the minimum faithful irreducible unitary representation of any compact connected Lie group acting on \mathbb{C}^2 .

This channel depends on two prior results: the 3-dimensional structure of the emerging space (Section 4.3, from the minimum H_1 condition), and the requirement that commitment amplitudes are complex (Section 10, from the Tsirelson derivation).

The space is \mathbb{C}^2 . By Theorem VIII, the minimum compact connected Lie group acting faithfully and irreducibly on \mathbb{C}^2 is SU(2). ✓

12.5 Channel 3: Extrinsic Curvature \rightarrow SU(3)

The third channel concerns the shape of the fold boundary — how it curves as one moves along it. This is extrinsic curvature: information about the embedding of the 2-dimensional surface in the ambient 3-dimensional structure.

For a 2-dimensional surface Σ embedded in 3-dimensional space, extrinsic curvature is encoded by the shape operator (second fundamental form) \mathbb{II} , a symmetric bilinear form on the 2-dimensional tangent space $T_p\Sigma$. In any orthonormal frame $\{e_1, e_2\}$ for $T_p\Sigma$:

$$\mathbb{II} = \begin{pmatrix} L & M \\ M & N \end{pmatrix}$$

Being a 2×2 real symmetric matrix, \mathbb{II} has exactly 3 independent real components: L, M, N . This is a theorem, not an assumption: a 2×2 symmetric matrix has $(2 \times 3)/2 = 3$ independent entries by construction.

Under complexification — required by Section 10's establishment that commitment amplitudes live in a complex inner-product space:

$$\mathbb{S}\mathrm{ym}^2(T^*\Sigma) \otimes_{\mathbb{R}} \mathbb{C} \cong \mathbb{C}^3$$

The space is \mathbb{C}^3 , provided the fold is not globally minimal on the active support of commitment conversion — which is established by the Proposition in Section 12.6.

By Theorem VIII, the minimum compact connected Lie group acting faithfully and irreducibly on \mathbb{C}^3 is $SU(3)$. ✓

12.6 Why the Boundary Cannot Be Minimal ($H \neq 0$)

The 3 independent components L, M, N would collapse to 2 if the fold boundary were a *minimal surface* — one with zero mean curvature $H = L + N = 0$. The trace condition would then fix one degree of freedom, leaving only $L - N$ and M independent, giving $\mathbb{R}^2 \rightarrow \mathbb{C}^2$ for the curvature channel. That would yield $SU(2)$ from the curvature channel rather than $SU(3)$, and the three-channel decomposition would produce only two distinct gauge sectors.

It is therefore essential to establish that the fold boundary cannot be globally minimal on the region where commitment conversion is active.

Proposition — Non-Minimality on Active Support Let Σ be a fold boundary carrying local commitment-conversion density $J(x)$. Suppose the effective fold action is:

$$\mathcal{S}\{\text{fold}\}[\Sigma] = \int_{\Sigma} \sigma, dA + \int_{\Sigma} U(\Phi, \mathcal{I}, J), dA$$
where σ is geometric tension and U encodes void-pressure Φ , informational loading \mathcal{S} , and commitment-conversion density J . The first variation gives an Euler–Lagrange equation of the form: $2\sigma H = J_{\text{eff}}$ where J_{eff} is the normal source term induced by U . Therefore on any open set where $J_{\text{eff}} \neq 0$, one has $H \neq 0$. An actively converting fold cannot be globally minimal; at most, source-free patches may satisfy $H = 0$.

The curvature channel therefore does not require $H \neq 0$ *everywhere* — it requires that the fold not be globally minimal on the active support of commitment conversion. Here "active support" means the subset of Σ on which commitment-conversion density J is nonzero — the region where irreversible fact formation is occurring. On source-free patches H may vanish; those patches contribute no curvature-channel information. On the active support — the region where the fold is genuinely converting void substrate into committed facts — $J_{\text{eff}} \neq 0$ and $H \neq 0$ follows from the Euler–Lagrange structure.

Status. The Proposition establishes a conditional chain: $H \neq 0$ on the active support follows from $J_{\text{eff}} \neq 0$, which follows from U being non-trivial, which requires explicit expressions for the void-pressure Φ , informational loading \mathcal{S} , and commitment-conversion density J derived from the void substrate. The Proposition as stated therefore establishes the *form* of the argument — that an active commitment interface cannot be globally minimal — rather than the claim itself. Three steps remain open: (i) deriving Φ, \mathcal{S} , and J explicitly from the void–fold axioms; (ii) showing that these produce a non-trivial U on the active support; and (iii) verifying that the resulting J_{eff} is generically nonzero. These are stated in Section 15. Until (i)–(iii) are complete, the $H \neq 0$ result is a structural argument whose conclusion is conditional on conditions that are themselves open.

12.7 Exhaustiveness and the Full Derivation

The three channels are exhaustive by the fundamental theorem of surface theory (Gauss-Codazzi equations): the first and second fundamental forms determine the local embedding of a surface in 3-space completely up to rigid motion. No fourth channel type exists.

The derivation in full, under the stated assumptions:

1. Axiom V1: the void admits no persistent asymmetry (Section 2)
2. The fold boundary Σ is the structural consequence of Axiom V1 (Section 3)
3. Σ is 2-dimensional — the minimum dimensionality for topological closure (Section 4.3)
4. The emerging space is 3-dimensional — treated here as a structural consequence of the coexistence of multiple independent closure regions (Section 4.3); a full proof of this dimensional claim remains to be given and should be regarded as an assumption at this stage
5. Commitment amplitudes inhabit a complex inner-product space (Assumption, Section 10)
6. By the fundamental theorem of surface theory, local geometric information at Σ decomposes exhaustively into scalar, spinorial, and extrinsic curvature channels
7. These channels require complex vector spaces of dimensions 1, 2, and 3 respectively under complexification
8. The extrinsic curvature channel requires that the fold not be globally minimal on the active support of commitment conversion; on that support $J_{\text{eff}} \neq 0$ and hence $H \neq 0$, so the shape operator retains 3 independent components (Proposition — Non-Minimality on Active Support, Section 12.6, status: structurally argued, requires explicit expressions for J_{eff})
9. By Theorem VIII: the minimum compact connected Lie groups acting faithfully on \mathbb{C}^1 , \mathbb{C}^2 , \mathbb{C}^3 are, respectively, $U(1)$, $SU(2)$, $SU(3)$. The three-channel decomposition selects $SU(3) \times SU(2) \times U(1)$ as the minimal compact Lie group structure compatible with the channel dimensions under the stated assumptions.
10. Their direct product is **$SU(3) \times SU(2) \times U(1)$**

12.8 Honest Assessment

The result is therefore a conditional structural derivation: if the fold boundary realises the stated geometric and amplitude assumptions, then the minimal compatible gauge-group structure is $SU(3) \times SU(2) \times U(1)$. The conditions and their status are assessed below.

An independent derivation of the same non-abelian sector appears in the companion paper *Structural Sufficiency and the Unique Non-Abelian Gauge Algebra*. There the three-dimensional excitation space of the fold is required to support complex irreducible transformations preserving interaction invariants. Under these conditions the Lie algebra $\mathfrak{su}(3)$ emerges as the unique non-abelian algebra acting faithfully on the three-dimensional sector. The appearance of $SU(3)$ in both the geometric boundary-channel argument developed here and the algebraic excitation argument of that work provides a useful consistency check across the VERSF program.

What is derived at the structural level: The existence of exactly three channel types from fold boundary geometry (standard differential geometry). The (1, 2, 3) complex dimensions from scalar, spinorial, and curvature structure respectively. The argument that active commitment

prevents the fold from being globally minimal on the active support, and therefore generically preserves a three-component curvature channel — though this requires explicit expressions for J_{eff} to formalise. $SU(3) \times SU(2) \times U(1)$ from (1, 2, 3) via Theorem VIII.

What remains to be formalised: The precise form of $\mathcal{S}_{\text{fold}}$ — the commitment density term \mathcal{S} and void-pressure Φ need explicit expressions derived from the void substrate. Showing formally that the Euler-Lagrange condition for $\mathcal{S}_{\text{fold}}$ has no solutions with $H = 0$ everywhere. Connecting the abstract channel structure to specific Standard Model particle representations, including hypercharge assignments. Most importantly: deriving the chirality of the $SU(2)$ coupling. Section 12 derives $SU(2)$ as the gauge group of the orientation channel, but does not establish that $SU(2)$ must couple exclusively to one Weyl sector. The result is $SU(2)$, not yet $SU(2)_L$. The chirality question is addressed separately in the VERSF channel-gauge correspondence program.

The empirical anchor question. The BCB framework (Bit Conservation and Balance), developed independently by the same author, derives the same gauge group from a different starting point: distinguishability conservation, entropy minimisation, and the Fisher information geometry of internal quantum state spaces. That program (Taylor, "Closing the Interfaces") identifies three minimal empirical anchors that it argues are genuinely required and cannot be eliminated by any purely structural argument: (A) stable three-body confined bound states exist, selecting $n = 3$ and $SU(3)$; (B) a fundamental two-state internal sector with continuous symmetry exists, selecting $n = 2$ and $SU(2)$; (C) quantum phases are unobservable (Born rule structure), selecting $U(1)$. The VERSF fold boundary derivation presented here arrives at the same (1, 2, 3) dimensions via geometric reasoning: the 3D embedding forces spinorial structure for \mathbb{C}^2 , and $H \neq 0$ is argued to force three independent shape operator components for \mathbb{C}^3 . Whether the VERSF geometric arguments genuinely replace BCB's Anchors A and B — making the derivation fully structural — or whether the same empirical content is present in different form (encoded, for example, in the assumption that the emerging space is 3-dimensional and the fold is genuinely active) is a question that deserves explicit treatment. The $H \neq 0$ argument bears the weight that Anchor A carries in BCB: both are required to prevent a collapse of the curvature channel from \mathbb{C}^3 to \mathbb{C}^2 . Whether one is strictly stronger than the other is not yet established.

Relationship to BCB. VERSF and BCB are parallel derivations of $SU(3) \times SU(2) \times U(1)$ from different pre-geometric starting points. BCB uses information-theoretic geometry (Fisher manifolds, distinguishability conservation, entropy minimisation). VERSF uses topological boundary geometry (void substrate, fold, commitment partial order). Both avoid assuming gauge structure. Both arrive at the same gauge group. Both identify the same open questions: gauge dynamics, coupling constants, symmetry breaking, and fermion generations are not derived by either framework. The two programs are complementary, not competing, and the consistency of their conclusions strengthens the case that $SU(3) \times SU(2) \times U(1)$ is not an arbitrary choice but a structural constraint imposed by the geometry of information at the boundary between reversible and irreversible structure.

Significance. Among geometric and topological approaches to fundamental physics, no prior framework derives $SU(3) \times SU(2) \times U(1)$ from the geometry of a pre-geometric boundary without assuming gauge structure, extra dimensions, or specific algebraic input. String theory produces a

landscape. Connes' noncommutative geometry recovers the gauge group from spectral data but requires assumed algebraic structure. The VERSF result is that the three-channel decomposition of an active 2D commitment interface in 3D space, combined with the assumption that commitment amplitudes are complex, selects $(\mathbb{C}^1, \mathbb{C}^2, \mathbb{C}^3)$ under the stated conditions and hence yields $SU(3) \times SU(2) \times U(1)$ via Theorem VIII. The gauge group is not imposed — it is what the local geometry of fact formation selects, subject to the caveats above regarding empirical content and the chirality gap.

13. The Synthesis

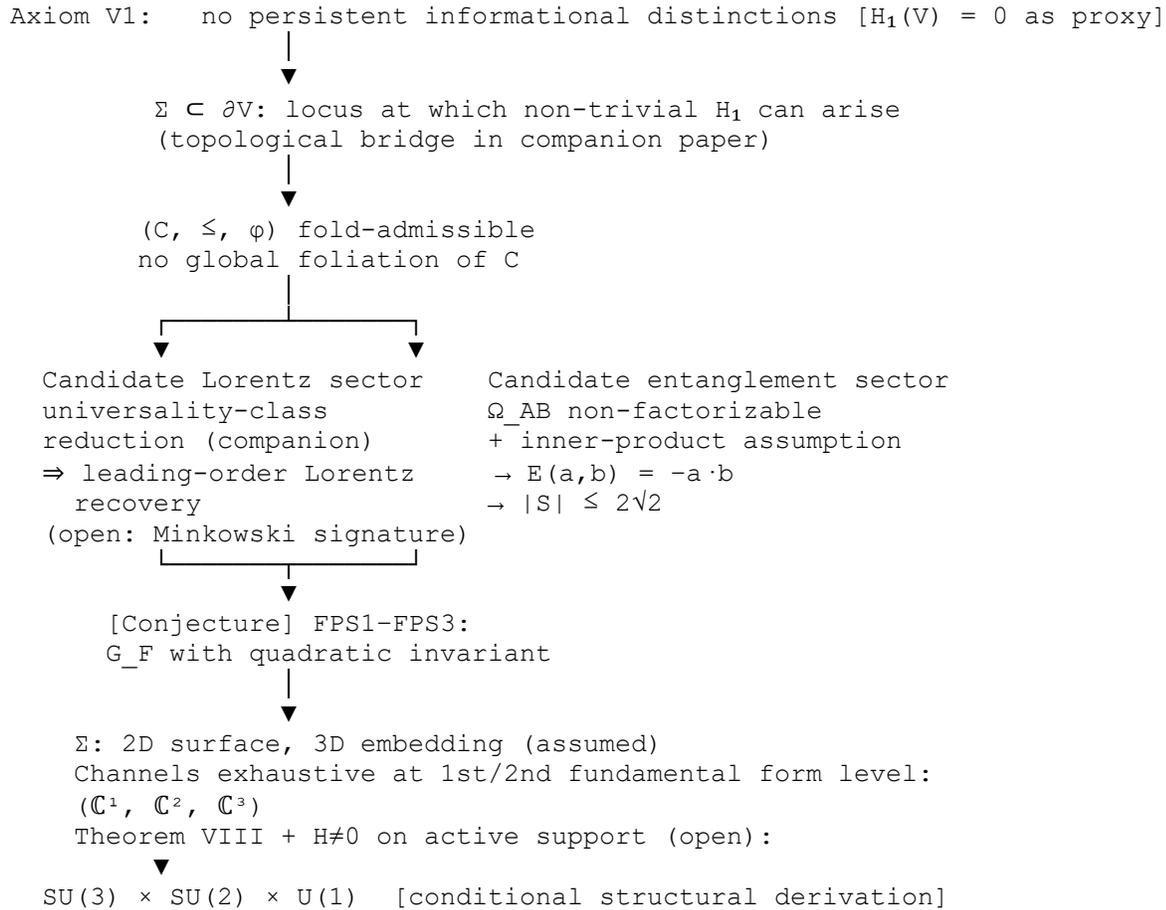
The eight layers of the framework — void, fold, topology of facts, commitment structure, Lorentz emergence, quantum correlations, fold projection symmetry, and Standard Model gauge group — are not independent. They are the same argument viewed from eight vantage points:

Layer	Core Claim	Formal Object
The Void	Substrate cannot sustain asymmetry	Zero-entropy boundary condition, $H_1(V) = 0$
The Fold	Shed asymmetry constitutes the boundary	Topological transition, minimal structure
Topology of Facts	Topological closure is the condition for facts	Non-trivial H_1 , isolated discard region D
Commitment Structure	Facts form a partial causal order	Fold-admissible triple (C, \leq, φ)
Lorentz Invariance	No foliation \rightarrow no preferred frame	Invariant relational statistics under coarse-graining
Quantum Correlations	Entanglement is a non-factorizable substrate relation	Joint commitment amplitude; Tsirelson bound $2\sqrt{2}$
Fold Projection Symmetry	Both constraints are one quadratic projection principle	G_F acting on H_F ; FPS1–FPS3; Unified Conjecture
Standard Model Gauge Group	Channel geometry at fold boundary \rightarrow $SU(3) \times SU(2) \times U(1)$	Three channels $(\mathbb{C}^1, \mathbb{C}^2, \mathbb{C}^3)$; Theorem VIII; $H \neq 0$

The logical chain is:

Axiom V1 (no persistent asymmetry in V) \rightarrow fold boundary Σ is structurally required \rightarrow topological closure at Σ \rightarrow first non-trivial H_1 \rightarrow commitment event \rightarrow fold-admissible partial causal order without global foliation \rightarrow (spacetime sector) no preferred frame is selectable \rightarrow Lorentz covariance under coarse-graining; (entanglement sector) non-factorizable joint substrate relation \rightarrow singlet correlation law \rightarrow Tsirelson bound $2\sqrt{2}$; (unification, conjectural) FPS1–FPS3 select both as representations of G_F preserving one quadratic invariant; (gauge structure) 2D fold boundary in 3D space \rightarrow three exhaustive geometric channels $(\mathbb{C}^1, \mathbb{C}^2, \mathbb{C}^3)$ \rightarrow active interface selects $H \neq 0$ (structurally motivated) \rightarrow by Theorem VIII: $U(1) \times SU(2) \times SU(3)$.

The same chain in compact mathematical notation:



Each step is necessitated by the previous. The gauge group derivation is the most concretely completed result; the FPS unification is the primary open conjecture; chirality and the BCB comparison are the two remaining gaps in the gauge sector.

The Tsirelson and gauge-group derivations share a common dependency on the complex inner-product structure introduced in Section 10. The gauge derivation additionally depends on the dimensional emergence argument of Section 4.3. These dependencies do not create circularity but indicate that several major results arise from a shared structural assumption. The origin of the inner-product structure within the void-fold axioms is therefore the deepest foundational open problem, noted first in Section 15.

14. Addressing Standard Objections

14.1 The Bootstrapping Problem

Objection: If the void is zero entropy, what initiates the entropy dynamics?

Response: Nothing initiates them. The fold is not an event; it is a structural necessity. The void's symmetry constraint implies a boundary, and that boundary is the fold. There is no prior moment at which the void "decides" to shed asymmetry — the fold is constitutive of what a zero-entropy substrate bounded by non-zero-entropy content means.

14.2 The Circularity of Time

Objection: Entropy change requires time; time cannot emerge from entropy change.

Response: Time in VERSF is not defined as a rate of entropy change but as the ordered accumulation of irreversible facts. Once a topological closure separates alternatives, that distinction becomes permanent, and the sequence of such closures defines temporal ordering. Time is not the parameter in which entropy changes — it is the index of fact formation. No circularity arises because time is downstream of the topological threshold, not presupposed by it. This aligns with the argument developed in the companion paper *The Topological Threshold for Fact Formation* that the arrow of time arises from the accumulation of irreversible distinctions rather than from a fundamental temporal parameter.

14.3 The Lattice Lorentz Objection

Objection: Discreteness implies a preferred frame.

Response: This objection presupposes a substrate with a global simultaneity structure that boosted observers can slice differently. VERSF's substrate has no such structure. Discreteness in VERSF is the discrete formation of commitment events, not the discrete spatial extension of a lattice. The commitment structure (C, \leq, ϕ) explicitly excludes any global foliation. There is no lattice to break.

14.4 The Bell Objection

Objection: Any discrete commitment-based framework assigns pre-existing values to particles. Bell's theorem then applies and the framework predicts $|S| \leq 2$, in contradiction with experiment.

Response: VERSF does not assign pre-existing local values. The void substrate does not factor an entangled pair into two independent hidden states. Before measurement, the pair is a single non-factorizable substrate relation — one distributed pre-geometric object, not two local ones. The joint commitment amplitude is explicitly non-Bell-local: $P(A, B | a, b, \Omega_{AB}) \neq \int d\lambda \rho(\lambda) P(A|a, \lambda) P(B|b, \lambda)$. Bell's theorem requires the factored form; VERSF rejects the factoring.

14.5 The FPS Import Objection

Objection: The Fold Projection Symmetry axioms — particularly FPS3's inner-product norm law — are imported from quantum mechanics rather than derived from the void substrate.

Response: This objection correctly identifies the residual gap stated in Section 10.5 and the derivation program of Section 11.6. FPS3 is currently adopted as a constrained admissibility

principle on fold projection geometry rather than a theorem derived from void topology alone: among projection-composition laws compatible with FPS1–FPS2, the quadratic norm law is the minimal one that introduces no additional relational structure. Whether this minimality condition follows strictly from Axiom V1 is the primary open problem. The paper does not claim FPS3 is derived — it claims it is the minimal admissible postulate and identifies what a full derivation would require. A referee who insists on derivation before publication is correct to do so; this is honestly presented as a constrained conjecture with a program.

14.6 The Gauge Group Circularity Objection

Objection: The derivation of $SU(3) \times SU(2) \times U(1)$ relies on Theorem VIII, which establishes that $(\mathbb{C}^1, \mathbb{C}^2, \mathbb{C}^3)$ uniquely determine $(U(1), SU(2), SU(3))$. But Theorem VIII is a mathematical result about Lie groups and their representations — it does not itself come from the void substrate. Isn't the gauge group derivation therefore only as good as Theorem VIII?

Response: Theorem VIII is a pure mathematical theorem about minimal compact connected Lie groups and faithful irreducible unitary representations. It makes no physical assumptions and requires no input from VERSF. It is not derived from the void — it is a result in representation theory that VERSF uses as a mathematical tool, in the same way that the paper uses differential geometry to characterise the shape operator. The physical content of the gauge group derivation lies entirely in establishing that the fold boundary produces channel spaces of dimensions exactly 1, 2, and 3. Theorem VIII then converts that geometric result into a group-theoretic result. The derivation is not circular: geometry \rightarrow channel dimensions \rightarrow (Theorem VIII) \rightarrow gauge groups.

14.7 The Uniqueness Problem

Objection: Why does this substrate give rise to *this* physics?

Response: The question has multiple parts. For the gauge group *structure*: the fold boundary geometry selects three channel types with complex dimensions 1, 2, and 3 as the minimal compatible decomposition, and Theorem VIII selects $SU(3) \times SU(2) \times U(1)$ as the minimal compact Lie group structure compatible with those dimensions. For the gauge group *chirality*: Section 12 derives $SU(2)$ but does not yet establish that it couples to only one Weyl sector. The result is $SU(2)$, not $SU(2)_L$. Chirality requires an additional argument — the BCB framework provides one (bidirectional coupling produces two independently conserved distinguishability currents, violating BC1's single conservation law) — but the corresponding VERSF argument remains to be developed. For the spacetime structure: fold admissibility constrains the class of emergent partial orders, and whether this selects Minkowski geometry as the unique solution is the primary open problem of Section 15. The framework makes substantial progress on the uniqueness question; it does not yet close it entirely.

14.8 The Empirical Anchor Objection

Objection: The BCB framework argues that three empirical anchors are irreducibly required to select the specific gauge factors: the existence of three-body confined bound states (for $SU(3)$), a two-state internal sector with continuous symmetry (for $SU(2)$), and the Born rule phase

structure (for $U(1)$). If those anchors are genuinely unavoidable, then the VERSF fold boundary derivation either implicitly encodes the same empirical content or overclaims.

Response: This is a legitimate and important challenge. The VERSF derivation and BCB derivation are parallel programs that should be evaluated together. Three points bear on the comparison.

First, the Born rule phase structure (Anchor C in BCB) corresponds in VERSF to the requirement that commitment amplitudes live in a complex inner-product space, established in Section 10 via the Tsirelson derivation. Both frameworks require this; neither derives it from more primitive principles. The empirical content is equivalent.

Second, the selection of $n = 3$ in BCB (Anchor A, three-body confinement) corresponds in VERSF to the extrinsic curvature channel requiring $H \neq 0$ to retain three independent shape operator components. The VERSF argument is structural — $H \neq 0$ follows from the fold being an active interface — but it is not yet a formal derivation, and whether it fully replaces the empirical content of Anchor A is not established. The honest position is that both frameworks require something here, and the VERSF version of that something is $H \neq 0$.

Third, BCB's Anchor B (two-state internal sector, selecting $n = 2$) corresponds in VERSF to the fold boundary being embedded in 3-dimensional emerging space, which forces spinorial character for the orientation channel. The 3D emergence is derived in Section 4.3 from the minimum H_i condition. Whether this constitutes a structural replacement for Anchor B or merely encodes it differently is a genuine open question.

The derivations converge on the same gauge group by different routes. Whether the VERSF route eliminates the need for the BCB anchors, or relocates their empirical content into geometric assumptions, requires explicit comparative analysis that has not yet been performed.

15. Open Problems

The framework raises the following questions, grouped by maturity. The first three constitute the derivation program for the Unified Fold Symmetry Conjecture. The fourth and fifth concern the gauge group derivation. The remainder are longer-range.

Origin of the inner-product structure. The complex inner-product geometry underlying both the Tsirelson bound and the curvature-channel complexification is currently inherited from the minimal quantum realization of commitment states (Section 10.4). Deriving this structure directly from the void-fold axioms would close the deepest remaining foundational gap in the framework.

The fold symmetry group. Define G_F , the fold symmetry group acting on the pre-geometric relational space H_F , and derive it from the void symmetry constraint and fold topology. Characterise which transformations of the fold leave observable commitment statistics invariant.

The two sector representations. Show that one representation of G_F generates the coarse-grained spacetime interval structure with Minkowski signature, and that a distinct representation generates the measurement-context overlap structure yielding $E(\mathbf{a}, \mathbf{b}) = -\mathbf{a} \cdot \mathbf{b}$ and the Tsirelson bound.

The single quadratic invariant. Show that there is a unique or natural quadratic form preserved by G_F , manifesting as Minkowski interval structure in the spacetime sector and as the bounded correlation kernel in the entanglement sector. These three together would establish the Unified Fold Symmetry Conjecture as a theorem.

Formalising the gauge group derivation. Derive the explicit form of $\mathcal{S}_{\text{fold}}$ — the commitment density \mathcal{S} and void-pressure Φ need precise expressions from the void substrate — and prove formally that the Euler-Lagrange condition for $\mathcal{S}_{\text{fold}}$ has no solutions with $H = 0$ everywhere. Establish the hypercharge assignments from anomaly cancellation in the VERSF context.

Chirality from fold geometry. Establish that the $SU(2)$ orientation channel couples to exactly one Weyl sector. The BCB framework proves this from distinguishability current conservation: bidirectional $SU(2)$ coupling produces two independently conserved internal currents, violating the single-conservation-law requirement of BC1. The corresponding VERSF argument — that the fold's commitment structure forces a single chirality — needs to be developed. Until this is done, the derivation yields $SU(2)$, not $SU(2)_L$.

Comparison with BCB. Establish explicitly whether the VERSF fold boundary argument eliminates the need for BCB's three empirical anchors or relocates their content into geometric assumptions about the fold. The $H \neq 0$ condition and Anchor A play structurally parallel roles; whether one is strictly stronger requires comparative analysis. If the two derivations can be shown to make identical assumptions in different languages, this would constitute independent confirmation that $SU(3) \times SU(2) \times U(1)$ is a structural consequence of the geometry of information at the reversible-irreversible boundary.

The admissibility class and Lorentzian universality. The Lorentz sector contains two independent open problems of different character and depth that should not be conflated.

(a) Universality-class selection. Does fold admissibility select partial orders that fall into the Lorentzian universality class established by the Reduction Result of Section 7.5? This is the more tractable of the two problems: the reduction result already shows that fold-compatible systems with a four-component internal fiber can exhibit Lorentz recovery at leading order. The remaining question is whether fold admissibility is sufficient to force membership in this class, or whether additional structure is needed.

(b) Minkowski signature. Even if fold-admissible partial orders fall into the right universality class, the coarse-grained bilinear form must have Minkowski rather than Euclidean signature $(+, -, -, -)$. This is the harder problem and deserves a more developed statement of what resolving it would require. It is not enough to show that the emergent geometry is Lorentz-covariant in the abstract; it must be Lorentzian specifically, with one timelike and three spacelike directions. The

mechanism proposed by VERSF is that fold-admissible partial orders have a fundamentally causal character — the dominant commitment chains are timelike rather than spacelike, so the long-range interval structure is hyperbolic rather than elliptic. This suggests that the partial order possesses an intrinsic arrow of precedence not symmetric under exchange of temporal and spatial directions, which is what would distinguish Minkowski from Euclidean signature at the coarse-grained level. Establishing this requires showing: (i) that the density of fold-admissible partial orders is anisotropic in a way that privileges causal chains; and (ii) that this anisotropy survives coarse-graining to produce a $(+, -, -, -)$ rather than $(-, -, -, -)$ or $(+, +, +, +)$ bilinear form. No current argument in the VERSF program addresses (i) or (ii) directly. Problem (a) is plausibly closable by extending the universality-class argument; problem (b) requires substantially new input about the causal geometry of fold-admissible orders.

The interface metric. What is the geometry of the fold boundary, and is it uniquely determined by the void's symmetry group or does a family of solutions exist?

The sequence of folds. If multiple folds occur, does their ordering carry physical information, and is this what gives rise to the full structure of spacetime?

16. Conclusion

The framework developed here consists of eight layers — void, fold, topology of facts, commitment structure, Lorentz emergence, quantum correlations, fold projection symmetry, and Standard Model gauge group — that are structurally connected rather than independently posited.

Under Axiom V1, the void substrate V admits no persistent distinguishable states ($H_1(V) = 0$). This requires a boundary Σ at which distinguishable structure accumulates. Topological closure at Σ is the condition for fact formation. Facts constitute commitment events that form a fold-admissible partial causal order without global foliation. The non-foliability of the commitment structure prevents any preferred inertial frame from being defined. The non-factorizability of joint commitment amplitudes yields singlet correlations bounded by the Tsirelson limit. Under the Fold Projection Symmetry axioms, both results are selectable as representations of a single fold symmetry group G_F — the Unified Fold Symmetry Conjecture. And the three-channel decomposition of Σ as an active 2D boundary in 3D space, under Theorem VIII, yields $SU(3) \times SU(2) \times U(1)$.

Three caveats govern the gauge group result. First, the $H \neq 0$ condition is structurally motivated but not yet formally derived from $\mathcal{S}_{\text{fold}}$. Second, the derivation selects $SU(2)$ rather than $SU(2)_L$ — chirality is not yet established. Third, the relationship between this geometric derivation and the BCB program's information-theoretic derivation of the same gauge group has not yet been formally compared.

Within this hierarchy, the gauge-group result is the most structurally developed, the Lorentz sector is reduced to two independent open problems — universality-class selection (Section 7.5)

and Minkowski signature emergence (Section 15) — and the FPS unification remains conjectural.

What the paper establishes is the logical architecture, with every assumption and open step explicitly identified. A single pre-geometric structure — the fold boundary Σ of a void substrate constrained by Axiom V1 — is sufficient, under stated assumptions, to organise a unified derivational program for the causal structure of spacetime, the quantum statistics of entangled systems, and the gauge-group structure of the Standard Model. The present work establishes the structural architecture of this program rather than its complete dynamical realisation. Whether these results follow from a single unified derivation rather than three parallel ones is the primary open question, stated precisely as the Unified Fold Symmetry Conjecture.

The VERSF program therefore consists of three principal remaining derivational tasks: (i) deriving the complex inner-product structure from Axiom V1 and the fold topology alone, without reliance on the quantum-substrate premise; (ii) deriving Minkowski rather than Euclidean signature from the causal geometry of fold-admissible orders; and (iii) deriving the chirality of the SU(2) coupling from fold boundary geometry. Completing these three tasks would convert the present conditional structural derivation into a closed derivation of spacetime geometry, quantum statistics, and gauge structure from a single pre-geometric axiom.

References

Causal set theory

- Bombelli, L., Lee, J., Meyer, D., Sorkin, R.D. (1987). Space-time as a causal set. *Physical Review Letters* 59, 521.
- Rideout, D.P., Sorkin, R.D. (2000). Classical sequential growth dynamics for causal sets. *Physical Review D* 61, 024002.
- Sorkin, R.D. (2003). Causal sets: Discrete gravity. *Lectures on Quantum Gravity*, Valdivia Summer School.

Bell inequalities and Tsirelson bound

- Bell, J.S. (1964). On the Einstein-Podolsky-Rosen paradox. *Physics* 1, 195.
- Cirel'son, B.S. (1980). Quantum generalizations of Bell's inequality. *Letters in Mathematical Physics* 4, 93.
- Tsirelson, B.S. (1987). Quantum analogues of the Bell inequalities: the case of two spatially separated domains. *Journal of Soviet Mathematics* 36, 557.
- Landau, L.J. (1987). On the violation of Bell's inequality in quantum theory. *Physics Letters A* 120, 54.
- Clauser, J.F., Horne, M.A., Shimony, A., Holt, R.A. (1969). Proposed experiment to test local hidden-variable theories. *Physical Review Letters* 23, 880.

Quantum information geometry

- Braunstein, S.L., Caves, C.M. (1994). Statistical distance and the geometry of quantum states. *Physical Review Letters* 72, 3439.
- Wootters, W.K. (1981). Statistical distance and Hilbert space. *Physical Review D* 23, 357.

Differential geometry

- do Carmo, M.P. (1976). *Differential Geometry of Curves and Surfaces*. Prentice Hall.
- Kobayashi, S., Nomizu, K. (1963/1969). *Foundations of Differential Geometry*, Vols. I & II. Wiley.

VERSF companion papers

- Taylor, K. The One Fold. VERSEF Theoretical Physics Program.
- Taylor, K. The Topological Threshold for Fact Formation. VERSEF Theoretical Physics Program.
- Taylor, K. Structural Conditions for the VERSEF Framework. VERSEF Theoretical Physics Program.
- Taylor, K. Structural Sufficiency and the Unique Non-Abelian Gauge Algebra. VERSEF Theoretical Physics Program.
- Taylor, K. Area Scaling of TPB Information Capacity. VERSEF Theoretical Physics Program.

Gauge group derivation (BCB program)

- Taylor, K. The Minimal Internal Symmetry Theorem in the BCB Framework. VERSEF Theoretical Physics Program.
- Taylor, K. Distinguishability Conservation and Gauge Structure. VERSEF Theoretical Physics Program.
- Taylor, K. Closing the Interfaces: Empirical Anchors, Abelian Uniqueness, and Chirality Selection in the BCB Framework. VERSEF Theoretical Physics Program.

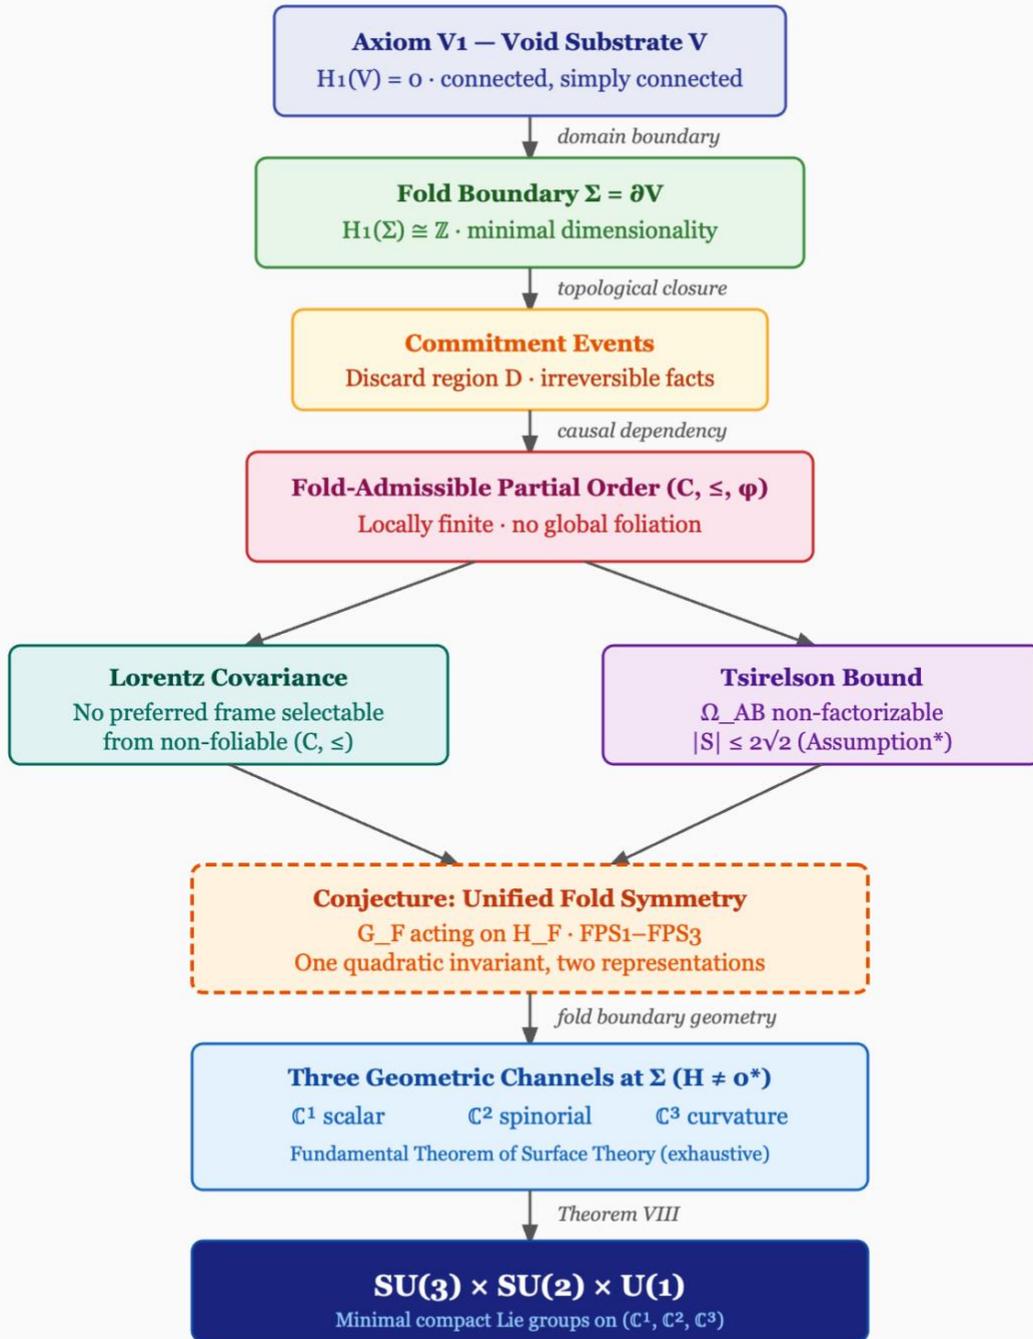
Emergent spacetime and noncommutative geometry

- Connes, A., Lott, J. (1991). Particle models and noncommutative geometry. *Nuclear Physics B (Proc. Suppl.)* 18, 29.
- Henson, J. (2006). The causal set approach to quantum gravity. In *Approaches to Quantum Gravity*, ed. Oriti. Cambridge.

Anomaly cancellation

- Minahan, J.A., Ramond, P., Warner, R.C. (1990). A comment on anomaly cancellation in the Standard Model. *Physical Review D* 41, 715.
- Babu, K.S., Mohapatra, R.N. (1990). Quantization of electric charge from anomaly constraints and a Majorana neutrino. *Physical Review D* 41, 271.

VERSF Framework: From Void to Standard Model



Open / Assumed (marked *)

- * FPS₃ (norm bound) postulated; derivation from fold topology is central open problem
- * H ≠ 0 structurally motivated; formal proof from S_{fold} pending
- * Complex inner-product assumed; chirality (SU(2)_L) not yet derived