VERSF Results Catalogue

Void Energy–Regulated Space Framework: A Structured Assessment of the Programme

---

Abstract

For the general reader.

This document is a structured inventory of the results produced by the Void Energy–Regulated Space Framework (VERSF), a theoretical physics research programme that attempts to derive the laws of nature — including quantum mechanics, spacetime geometry, and the particle content of the Standard Model — from a small number of operational principles about how stable physical facts can exist at all.

The programme begins with a simple observation: for physics to be possible, a universe must support facts that are stable, recoverable, and distinguishable from one another. From this starting point, it derives three necessary conditions — facts must have a minimum size, they must be irreversible once formed, and a bounded region can only hold a finite number of them. Everything else in the framework follows from these three conditions and the structure of the boundary at which reversible possibilities become irreversible physical records.

The most striking results the programme claims are: that time emerges from the accumulation of irreversible facts rather than being a background dimension; that the quantum mechanical limit on correlations between distant particles (the Tsirelson bound) follows from the geometry of that boundary; that the symmetry groups of the Standard Model of particle physics — governing the electromagnetic, weak, and strong forces — arise from the three independent geometric channels available at a two-dimensional boundary embedded in three-dimensional space; and that there are exactly three generations of matter particles for structural rather than accidental reasons.

This document classifies these results honestly. Some are established as theorems within the programme's own axiomatic system. Others are conditional on additional assumptions that are themselves open to scrutiny. Others are physical identifications — places where a mathematical structure is matched to a measured quantity — that require further bridge steps not yet completed. The document also identifies failure conditions: what would count as falsification of specific sectors of the programme.

The document is long because the programme is large. Readers seeking an orientation should begin with the Dependency Spine (which shows how results build on each other), the Executive Catalogue (the ten headline results), the Status Table (a one-page overview of every result and its current state), and the Condensed Failure Conditions (what would count as failure). These appear in the front matter before the detailed sections.

Table of Contents

Programme Orientation

What This Document Is

The Void Energy–Regulated Space Framework (VERSF) is a multi-paper research programme that attempts to derive the structure of physical law from operational constraints on distinguishability, information, and irreversible record formation. Rather than taking spacetime or quantum fields as primitive, VERSF begins with the requirements for physical facts to exist at all. The present catalogue summarises results drawn from more than one hundred companion papers. Its purpose is not to present new derivations but to organise and classify results already developed across the VERSF corpus.

The programme is built on a layered architecture in which reversible distinguishability relations exist in a substrate called the **Void**; irreversible commitments create **facts**; those commitments accumulate to generate **time**; and the relational structure of commitments produces **spacetime geometry** and **physical law**.

This document organises results into three tiers: (I) theorems established within the VERSF axiomatic framework, (II) conditional derivations and explicit realisations dependent on additional structural inputs, and (III) failure conditions and falsification tests. This classification follows the standard adopted in the programme's own companion audit paper (*Old-Fold-Companion 5*), which explicitly separates unconditional theorems, conditional derivations, and physical identifications.

Scope: This catalogue does not claim that listed results are accepted by mainstream physics or have been externally peer-reviewed as established science. It claims they are the results presently established, conditionally derived, or proposed *within the VERSF research programme*, assessed honestly against the programme's own axioms. Part III identifies what would count as failure and the experimental tests that can confirm or rule out specific predictions.

Epistemic Boundary

VERSF does not assume that the Void's internal structure is directly describable using the categories of emergent physics. Since spacetime, measurement, and stable distinctions arise only after commitment, those concepts may not apply inside the Void in any ordinary sense. Applying them there would be a category error.

The programme therefore adopts an **interface-based methodology**: it studies only those structural properties of the Void that are inferable from the fold boundary, where reversible distinguishability becomes irreversible physical record. All downstream results are derived from this interface and its consequences. The fold is the real object of theory; the Void is characterised solely by what it cannot do — sustain persistent distinctions — and by the constraints this imposes at its boundary.

Dependency Spine of the Programme

The following chain shows how the programme's core results build on one another. Every result downstream depends on those upstream; a failure at any node propagates forward.

```
Pre-physical substrate (Void) — characterised only by absence of persistent
distinctions
        ↓ [epistemic boundary — interface only]
Unavoidable structural conditions for physics
  (A1 finite distinguishability,
   A2 irreversible commitment,
   A3 finite localisation capacity)
        ↓
Fold / commitment boundary
        ↓
Bit Conservation and Balance (BCB)
        ↓
Ticks-per-Bit (TPB)
        ↓
Admissibility (law-filter)
        ↓
Fact-production mechanism and rate equation
        ↓
Time as accumulated commitment
        ↓
Causal partial order
        ↙                    ↘
Spacetime emergence       Quantum structure
        ↓                         ↓
Lorentz recovery          Tsirelson bound
        ↓                         ↓
Role-4 field equations    Gauge group derivation
        ↓                         ↓
GR limit                  Particle closure ontology
                                  ↓
                          Fermion generations
                                  ↓
                          Mass hierarchy (FFP)
```

A failure to close any derivational gap — for example, deriving the commitment barrier Φc or the complex inner-product structure from the fold/interface axioms — does not collapse the entire programme, but it does leave the downstream results hanging from an ungrounded node.

Three-Layer Structure and Dependency Risks

The programme resolves into three layers: **Layer 1 (Foundations)** — finite distinguishability, irreversible commitment, finite localisation capacity, BCB, TPB, admissibility; **Layer 2 (Emergent structure)** — time, causal order, Lorentz recovery, quantum amplitude structure, Tsirelson bound; **Layer 3 (Physical identification)** — gauge group, particle closures, fermion generations, coupling constants, mass spectrum. Layer 1 results are the internal theorems of Part I; Layer 2 and 3 results are the conditional derivations and physical identifications of Part II.

Because the programme is strictly layered, failure propagates upward. The critical load-bearing node is the fold/interface: if the interface law, commitment barrier, and projection rules cannot be derived, the downstream spacetime, quantum, and particle results lose their grounding. Key remaining derivational targets include: the interface/fold law, the commitment barrier Φ_c, complex inner-product from fold axioms, chirality ($SU(2)_l$), and FFP mass numerics. The Born rule derivation has been completed via the Double Square Rule (§6.3).

Condensed Failure Conditions

The following are the conditions under which specific sectors fail. Full failure conditions appear in §13.

Sector	Failure condition
Substrate programme	Interface law and fold projection rules cannot be derived
Fact production	Φ_c must be fitted rather than derived
Quantum structure	Complex inner-product cannot be derived from fold axioms
Lorentz invariance	Fold-admissible structures found with non-Lorentzian coarse-graining
Gauge group	Chirality ($SU(2)_l$) cannot be established from fold geometry
Fine-structure constant	RG bridge requires a fitted intermediate parameter
Fermion masses	FFP numerics produce mass ratios outside stated uncertainties
Experimental	Laboratory predictions excluded at relevant precision

---

Executive Catalogue of Major Achievements

The following results represent the programme's headline contributions, grouped by character. Each is discussed in full in the sections below; this list is provided so that readers can orient themselves before engaging the detail.

Foundational achievements *(established within the VERSF axiomatic framework)*

1. **Unavoidable conditions for physics** — Three necessary structural conditions (finite distinguishability, irreversible commitment, finite localisation capacity) are derived as prerequisites for any universe capable of supporting stable physical laws accessible to finite observers. *(§1.1–1.4)*
2. **Time as accumulated commitment** — Time is defined and derived operationally as the ordered accumulation of irreversible commitment events, explaining temporal asymmetry without a background time dimension. *(§1.8)*
3. **Fact-production mechanism and rate equation** — The three-condition commitment criterion unifies quantum measurement and entropy production as instances of a single

nucleation-like process within a reversible substrate, with the fact-production rate equation as its quantitative expression. *(§2.1, §6.1)*

Structural physics results *(established or leading-order conditional within the programme)*

4. **Quartic capacity law** — Structural theorem: $\chi(L) = \rho L^4/\hbar c$ is the natural information-capacity invariant of causal regions, appearing independently in quantum speed limits, entropy bounds, and causal-diamond action calculations. Physical identification (separate step): applied to the observed vacuum energy density, this yields a coherence scale $\xi \approx 8 \times 10^{-5}$ m. *(§1.5)*
5. **Particle ontology as closure defects** — Particles arise as stable topological closure defects in the committed distinguishability graph, providing a unified ontology for matter derived from the same structure as spacetime. *(§5.1)*
6. **Tsirelson bound from fold geometry** — The maximum quantum correlation strength $|S| \leq 2\sqrt{2}$ is derived from structural constraints on fold-boundary projection symmetry, connecting quantum nonlocality to the same geometry governing spacetime emergence. *(§8.1)*
7. **Lorentz recovery at leading order** — Lorentz-covariant dispersion is recovered from the fold substrate at leading order, with violations suppressed by the discreteness scale. The non-foliable commitment partial order eliminates any preferred frame. *(§7.2–7.3)*

Conditional derivation programme *(conditional within the VERSF structural programme)*

8. **Standard Model gauge group** — The three independent geometric channels of the 2D fold boundary yield U(1), SU(2), SU(3) as minimal compact symmetry groups, producing SU(3) × SU(2) × U(1) structurally. *(§9.1)*
9. **Fine-structure constant** — A geometric coupling invariant $\alpha(\text{geom}) = 1/144$ is derived from fold closure structure; with explicit bridge assumptions this maps to $\alpha^{-1} \approx 137$. The structural and identification steps are kept strictly distinct. *(§9.2–9.3)*
10. **Three fermion generations** — The Fermion Fold Principle identifies three, and only three, stable topological minimisers of a Fisher-distinguishability functional, giving a structural rather than parametric origin for the observed three fermion generations. *(§11.4)*

Status Table

A navigational overview of the full programme. Detailed entries appear in Parts I–III.

Result	Tier	Key open element
Unavoidable conditions A1–A3	Part I — derived from axioms	—
Time as accumulated commitment	Part I — derived from axioms	—

Result	Tier	Key open element
Causal partial order	Part I — derived from axioms	—
No preferred frame (non-foliability)	Part I — derived from axioms	—
Four-dimensional internal Hilbert space	Part I — structurally motivated	Why this realisation is canonical (§4.1)
Particle identity from fibre uniqueness	Part I — structurally motivated	Correspondence to physical particles is Class III
Particle ontology as closure defects	Part I — structurally motivated	Correspondence to physical particles is Class III
Quartic capacity law $\chi = \rho L^4/\hbar c$	Part I — derived; ρ is a free parameter	Physical identification of ξ requires vacuum energy density input
Fact-production rate equation	Conditional (Part II)	Φc not derived from interface structure
Lorentz invariance at leading order	Conditional (Part II)	Generic universality class unproved
Tsirelson bound from fold geometry	Conditional (Part II)	Complex inner-product not derived from fold axioms
Standard Model gauge group	Conditional (Part II)	Channel decomposition needs explicit justification; chirality open
Geometric coupling invariant α(geom)	Conditional (Part II)	Democratic allocation needs independent motivation
Fine-structure constant $\alpha^{-1} \approx 137$	Physical identification (Class III)	Three bridge steps unfinished; circularity risk if allocation is reverse-engineered
Cosmological constant scaling	Conditional + Class III identification	Numerical match is identification, not prediction
Three fermion generations	Conditional (Part II)	FFP axioms; identification is Class III
Rigidity of Yukawa sector	Conditional (Part II)	Exact masses require numerical solution
Exact fermion masses	Open numerical problem	FFP fold equations unsolved
Interface/fold law	Open foundational problem	Projection rules and admissible states unspecified
Complex inner-product from fold axioms	Open foundational problem	Required for Tsirelson and gauge elevation to Part I
Chirality ($SU(2)_l$)	Open derivation	Required for complete Standard Model
Φc from interface structure	Open derivation	Required for quantitative rate predictions

Result	Tier	Key open element
Geometry-dependent decoherence	Experimental (untested)	Falsifiable by spin-coherence experiments
Sub-Tsirelson multipartite reduction	Experimental (untested)	Falsifiable by large-separation CHSH tests
Coherence-scale Casimir signal	Experimental (untested)	Falsifiable by precision Casimir measurements
CMB / early-galaxy signatures	Indirect cosmological test	Next-generation survey sensitivity required
GW propagation in Role-4 regime	Indirect astrophysical test	Third-generation detector sensitivity required

---

Interpretive Rule for the Catalogue

A result may be mathematically strong while its physical identification remains conditional. Throughout this catalogue, derivation and identification are treated as distinct achievements. The programme does not claim that every mathematical structure has already been uniquely matched to observed physics; it claims that a large subset of physical structures arise naturally from the fold/interface framework, with varying degrees of dependence on additional assumptions. When a result is listed as "established within the VERSF axiomatic framework," this means the derivation holds within that axiomatic system. When a result is listed as "conditional within the VERSF structural programme," the derivation requires additional inputs that are themselves open to scrutiny. When a result is a "physical identification," a separately motivated step is required to connect the mathematical object to the measured quantity. These three categories are always kept distinct.

Part I — Established Within the VERSF Axiomatic Framework

Results in this tier follow from the programme's base axioms without additional structural assumptions. "Established within the VERSF axiomatic framework" means established as theorems within that axiomatic system — it does not assert independent verification or acceptance by the broader physics community. Internal theorems are mathematical consequences of the VERSF axioms; their correspondence to empirical physics requires separate identification steps described in Part II.

This tier contains two sub-types, marked in each entry:

- *(Derived)* — the result follows by explicit argument from A1–A3 or from structural axioms with no additional assumed input.

- *(Structurally motivated)* — the result follows given a representational or ontological assumption that is itself strongly motivated by the framework but not yet derived from the base axioms alone. Sceptics should give these entries closer scrutiny.

---

1. Foundational Architecture and Ontology

1.1 Physics requires fact-stable state classes *(Derived)* Physics is possible only if a universe supports stable, recoverable, reproducible facts. This is the single primitive from which later axioms are derived — not a convenient starting point but a necessary precondition for any operational physical theory.

1.2 Finite distinguishability is necessary — Axiom A1 *(Derived)* Any fact-supporting universe must have a minimum resolvable granularity. Arbitrarily precise distinctions cannot be stably recorded by finite observers within causally bounded regions; stable physical facts therefore require a lower bound on distinction size.

1.3 Irreversible commitment is necessary — Axiom A2 *(Derived)* Fact formation requires irreversible export of correlations beyond local causal control. This is not an optional measurement postulate but a structural consequence of finite observers combined with causal boundedness. A distinction that can be undone is not yet a fact.

1.4 Finite localisation capacity is necessary — Axiom A3 *(Derived)* Once A1 and A2 hold, a finite upper bound on stabilisable record density follows. A bounded region cannot support an unlimited density of permanent, distinguishable records. Axioms A1–A3 together constitute the *unavoidable structural conditions* on any operational physical theory.

1.5 Quartic capacity of causal regions (Derived)

- *Structural theorem:* A fundamental capacity parameter governs bounded spacetime regions:

$$\chi(L) = \rho L^4 / \hbar c$$

Here ρ is a generic energy density — a free parameter in the structural result. The formula is not dimensional analysis: it arises from the combination of quantum speed limits, entropy bounds, and causal-diamond action in a specific way, and the same ρL^4 combination appears independently across all three contexts. The threshold $\chi \sim 1$ marks the minimum region capable of producing an irreversible physical record. *Note on scope:* the structural theorem holds for arbitrary ρ; fixing ρ to the vacuum energy density is the physical identification step below, which imports empirical input.

The appearance of the same ρL^4 combination across quantum speed limits, thermodynamic entropy bounds, and causal-diamond action calculations suggests that $\chi(L)$ is the natural information-capacity invariant of causal regions.

- *Physical identification (separate step):* When applied to the observed vacuum energy density, the quartic capacity relation yields a characteristic coherence scale:

$$\xi \approx 8 \times 10^{-5}\ \mathrm{m}$$

This is interpreted as the smallest region capable of sustaining a classical fact, and reappears independently in several otherwise unrelated parts of the programme. This identification is Class III: the structural theorem is internal; its mapping to the vacuum energy density is a separately motivated step (see §9.4).

1.6 The Void as reversible substrate *(Derived — architectural primitive)* The Void is the domain in which no persistent informational distinctions exist. Distinguishable structures cannot remain there and accumulate at a structural boundary called the *fold*, which is the geometric locus at which facts form.

1.7 Commitment events as primitive facts of history (Derived) Commitment events are the elementary transitions by which reversible distinctions become physical facts. The universe is an expanding ledger of commitments, ordered by their causal relations.

1.8 Time as accumulated commitment (Derived) VERSF time is defined operationally as the ordered accumulation of irreversible commitment events:

$$\tau = \int d\sigma_{irr}$$

Time is not a background parameter but an emergent ordering relation generated by fact formation. This provides a structural explanation for temporal asymmetry and its connection to entropy, without presupposing either.

1.9 Bit Conservation and Balance (BCB) *(Derived — with caveat)* The qualitative principle that a bounded region cannot sustain an unlimited density of stable committed records follows directly from A1–A3 by explicit argument: A3 asserts finite localisation capacity, and BCB is the bookkeeping formalisation of that constraint. The *qualitative* upper bound on record density is therefore genuinely derived. However, the *quantitative* BCB bound — the specific numerical capacity N(BCB) appearing in the rate law — requires additional structural input (the energy-capacity relationship, TPB) and is not derived from A1–A3 alone. If BCB is being claimed in the strong quantitative sense, additional inputs should be acknowledged. In the weaker qualitative sense, the (Derived) label is appropriate.

1.10 Ticks-per-Bit (TPB) as commitment-throughput structure *(Derived — with caveat)* The qualitative result that commitment throughput is bounded jointly by energy and available capacity follows from BCB and the energy-finiteness assumption inherent in A1–A3. BCB and TPB together define the quantitative skeleton of fact production. The same caveat applies as for BCB: the *structural* rate-bounding principle is derived; the *specific quantitative rate structure* entering the rate law (§6.1) requires additional inputs. If TPB is earning the (Derived) label for the structural rate-bounding result, that is justified. If it is being claimed for the full quantitative rate structure, additional justification is needed and the label should be (Structurally motivated).

1.11 Admissibility as the law-filter *(Derived — with caveat)* The *principle* that some admissibility filter must exist — that not all transitions can become committed records if A1–A3 hold — follows directly from the axioms: finite distinguishability, irreversible commitment, and finite localisation capacity together imply that only a restricted class of transitions can stably propagate as committed facts. Physical law, in this framework, is reframed as a selection rule on commitments.

Important caveat: The *specific content* of the admissibility filter — which transitions are admissible and why — depends on the interface law (an open problem), since admissibility is precisely the set of transitions that satisfy the fold boundary constraints. If the interface law is the most critical unresolved open problem, then admissibility in its full specific form cannot be derived from A1–A3 alone. The same qualitative/quantitative distinction applied to BCB and TPB applies here: the existence of an admissibility filter is derived; the exact filter content is not, and depends on the fold boundary structure that remains to be fully characterised.

---

2. Fact Production, Entropy, and Measurement

2.1 Three-condition commitment criterion (Derived) A commitment event occurs when three conditions are satisfied simultaneously: (i) the physical distinction exceeds a minimum distinguishability threshold; (ii) it is amplified into environmental degrees of freedom beyond causal recovery; (iii) sufficient capacity exists to store the resulting record. This criterion unifies quantum measurement and entropy production as instances of a single mechanism.

2.2 Entanglement Ledger (Derived) The Entanglement Ledger is the cumulative bookkeeping device tracking correlation export across history. It links commitment, decoherence, and irreversibility within a single accounting structure.

2.3 Measurement as fact production (Derived) Within the framework, quantum measurement is recast as fact production rather than treated as a separate primitive postulate. The measurement problem is dissolved by identifying it with the general question of when a commitment event occurs, governed by the three-condition criterion. Within the framework the Born rule appears as the statistical law governing commitment selection; its derivation from fold axioms is a conditional result treated in Part II §6.3.

The three-condition criterion answers *when* a commitment event occurs (dissolving the trigger question); *which* outcome occurs at what probability is governed by the Born rule, whose derivation from fold axioms is treated in §6.3. The Double Square Rule derivation establishes that the quadratic probability law is the unique assignment compatible with the programme's structural constraints, completing the dissolution of both the trigger and the probability questions.

---

3. Spacetime Structure

3.1 Topology of facts requires closure *(Derived)* Irreversible facts require nontrivial first-homology structure. The fold creates the minimal topology needed for permanent distinctions: one-dimensional structures cannot support the required topological closure; two dimensions are the minimum for a single irreversible fact to exist.

Floor vs ceiling clarification: This result establishes the *floor* — the minimum dimension in which a single stable closure can exist. It does not establish that 2D is sufficient for a physical universe; that would require only one closure to coexist with nothing else. The question of why physical space has three dimensions rather than two is answered by §7.1, which addresses the *relational* minimum: the minimum dimension in which multiple distinct stable closures can coexist and interact. §3.1 gives the floor (2D); §7.1 gives the argument for the ceiling in the sense of the relational minimum (3D). The 3D result is conditional (§7.1 is a sketch awaiting full proof); the 2D floor result is derived here. A reader wondering "if 2D is the minimum, why isn't the answer 2D?" should consult §7.1, which addresses exactly that question.

3.2 Commitment events form a causal partial order (Derived) The commitment structure is formalised as a partially ordered set $(C, \leq, \varphi)$ with local finiteness and no global foliation. This is the pre-geometric backbone from which spacetime emerges.

3.3 No global foliation, hence no preferred "now" (Derived) The commitment structure is non-foliable by construction. No preferred inertial slicing exists in the commitment partial order. This eliminates the standard objection that any discrete substrate selects a preferred frame at the level of the underlying structure.

---

4. Quantum Foundations

4.1 Four-dimensional internal Hilbert space *(Structurally motivated)* A commitment event carrying one binary distinction plus reversible $\mathbb{Z}_2$ directionality has minimal reversible quantum realisation $\mathscr{H}(\text{fold}) \cong \mathbb{C}^4$, with rays in $\mathbb{CP}^3$. Listed among bedrock results in the One-Fold robustness audit. *Open question:* the argument identifies the minimal quantum realisation given the assumption that a binary distinction with $\mathbb{Z}_2$ directionality is the right input structure — but why this is the canonical starting realisation rather than some other is not yet derived from A1–A3 alone. This entry sits closer to the structurally motivated sub-type than the derived sub-type.

4.2 Binary directionality *(Structurally motivated)* The $\mathbb{Z}_2$ directionality of the commitment boundary is strongly motivated by fold geometry, but its derivation from the base axioms without additional representational input remains incomplete.

4.3 Particle identity from fibre uniqueness *(Structurally motivated)* Within the framework, same-type particles share the same internal fibre structure. This is structurally well-motivated by the closure ontology; its correspondence to the empirically observed identity of elementary particles is a physical identification that depends on that ontology being correct.

5. Particle Ontology

5.1 Particles as stable closure defects *(Structurally motivated)* A particle is a stable, localised, transportable closure defect in the committed distinguishability graph Σ = (F, E). Particles emerge from the chain — reversible distinguishability → commitment at the fold → committed graph → stable closure → observed particle — rather than being inserted by hand. This is a structurally motivated ontological reinterpretation; its correspondence to physical particles is a conditional identification that depends on closure dynamics being correctly specified.

5.2 Matter particles versus gauge bosons *(Structurally motivated)* Matter particles are stable closures; gauge bosons are closure-preserving propagation modes rather than stable closure defects. This ontological distinction is strongly motivated by the closure structure, though its full derivation depends on the closure dynamics of §10.

Part II — Conditional Within the VERSF Structural Programme

Results in this tier are genuine technical achievements of the programme but carry dependence on additional structural inputs beyond the base axioms. "Conditional within the VERSF structural programme" means the result follows given the stated additional assumptions, which are themselves open to scrutiny. Each entry states its conditional dependence explicitly. Where a result is further split into a structural component and a physical identification, both steps are listed separately.

---

6. Fact Production Dynamics

6.1 Fact-production rate equation The programme proposes the core rate law:

$$\Gamma(\text{fact}) = R(\text{TPB}) \cdot \exp(-\Phi c/\Phi) \cdot (1 - N(\tau)/N(\text{BCB}))$$

This interprets fact formation as a nucleation-like process in a reversible substrate, unifying quantum measurement, entropy production, and the arrow of time.

- *Structural result:* Given the BCB/TPB substrate and the nucleation analogy, the functional form of the rate law is derived.
- *Open element:* The commitment barrier Φc is not yet derived microscopically from the BCB/TPB substrate . Until it is, the rate law has a structurally unmotivated free parameter.

6.2 Time-flow equation from fact production

dτ/dt = Γ(fact)

*Conditional within the VERSF structural programme — * the rate law (§6.1), and therefore inherits its open element.

6.3 Born rule derivation (Double Square Rule) The Double Square Rule establishes that the quadratic probability law $P = |\psi|^2$ is the unique probability assignment compatible with the programme's structural constraints. The derivation proceeds from nine axioms — including positivity, normalisation, relabelling invariance, factorisation on product systems, interference structure, and TPB-consistency — and shows that any alternative (higher-order kernel, non-bilinear form) produces a concrete failure: either normalization leakage under unitary evolution, gauge violation, or loss of distinguishability. The TPB-consistency argument is particularly decisive: the 2^(1−p/2) calculation under Hadamard evolution shows that normalisation is preserved under reversible commitment transformations if and only if $p = 2$. A second independent route via symmetry (rotational invariance of the Void substrate plus operational stability) also converges on the quadratic form.

This is not an interpretive reframing of an assumed Born rule; it is a derivation. The Born rule emerges as the unique outcome-probability law compatible with the fold commitment structure.

Conditional within the VERSF structural programme: the nine Double Square axioms and the TPB-consistency framework.

---

7. Spacetime Emergence and Lorentz Structure

7.1 Three-dimensional space from relational closure The programme argues that two spatial dimensions are the minimum for topological closure (§3.1), and that three dimensions emerge relationally as the minimum needed for stable closures to coexist and interact with one another. The argument sketch is: a single 2D closure can exist in isolation, but two closures interacting requires them to be embedded in a space with enough room to support their boundaries and intersection structure; the minimum dimension that admits two distinct stable 2D closures with non-trivial boundary relations is 3. Space with dimension 2 would force all closure boundaries to coincide or cross in ways that destroy independent stability; dimension 4 or higher permits stable coexistence but is not required for it. Three is therefore the minimum relational dimension.

This argument is a sketch, not a proof. The full generic proof — that this relational minimum is necessarily 3 for all fold-admissible closure configurations, not just specific examples — is acknowledged as incomplete. Until that proof exists, this result is a strongly-motivated conditional, not an established theorem.

Conditional within the VERSF structural programme: the 2D closure necessity theorem (§3.1) and the relational minimum-dimension argument. The full generic proof of the 2D → 3D step remains open .

7.2 Lorentz invariance at leading order A local Hermitian translation-invariant Hamiltonian with the minimal four-component fold fibre recovers Lorentz-invariant dispersion at leading order, with violations suppressed by the discreteness scale.

*Conditional within the VERSF structural programme — * fold-admissibility constraints. The full theorem — that all fold-admissible commitment structures necessarily flow to the Lorentz universality class — is open .

7.3 Lorentz interval statistics from commitment counting When the commitment partial order is coarse-grained, the number of commitment events between two loci approximates the volume of a causal diamond, reproducing Lorentzian interval statistics without a preferred inertial frame.

*Conditional within the VERSF structural programme — * coarse-graining procedure and the commitment partial order structure.

---

8. Quantum Foundations: Conditional Results

8.1 Tsirelson bound from fold geometry

- *Structural result:* Five structural constraints on joint commitment amplitudes — unbiased marginals, rotational invariance of the Void substrate, and three further geometric constraints — yield the singlet correlation law and the bound $|S| \leq 2\sqrt{2}$. A second independent derivation via entropy-curvature constraints on $\mathbb{C}P^3$ provides corroborating support.
- *Conditional within the VERSF structural programme — * the assumption that commitment amplitudes inhabit a complex inner-product space, which is inherited from fold-realisation analysis rather than derived directly from fold/interface axioms . Closing this gap would elevate the result to Part I.

8.2 Fold Projection Symmetry principle Both Lorentz covariance and the Tsirelson bound are organised by the same projection constraints; both sectors depend only on invariant quadratic overlaps of fold data. This is a unifying conjecture connecting quantum nonlocality and spacetime structure through a common geometric mechanism.

Status: Conjectural principle, not yet a theorem.

8.3 Uniqueness claims The programme argues that certain structures are uniquely forced within the framework's stated axioms. These claims vary significantly in strength and should not be treated uniformly.

- *Near-theorem status within the axiom set:* The Tsirelson bound at 2√2 given the five geometric constraints (§8.1) — the derivation is explicit and the constraints are clearly stated. The complex Hilbert space structure, given standard quantum mechanics as background, is also well-supported.
- *Structurally motivated but not yet theorem-level:* Three fermion generations from the FFP — this requires that the FFP axiom set itself be uniquely forced from fold geometry, which is not yet established; the uniqueness is within the FFP axioms, not derived from them. Similarly for the Yukawa scale.
- *Established result:* Uniqueness of the Born rule is established by the Double Square Rule derivation (§6.3), which shows the quadratic form is the unique assignment compatible with the programme's structural constraints.

The phrase "uniqueness within stated axioms" is the correct framing for all of these — but what that means varies from near-theorem to open conjecture depending on how well-grounded the stated axioms are.

---

9. Gauge Structure

9.1 Standard Model gauge group from fold-boundary geometry The 2D fold boundary embedded in emerging 3D space is argued to have exactly three independent local geometric channels:

Geometric channel	Complex space	Minimal compact connected faithful group
Scalar separation	$\mathbb{C}^1$	U(1)
Spinorial orientation	$\mathbb{C}^2$	SU(2)
Extrinsic curvature	$\mathbb{C}^3$	SU(3)

The identification of minimal compact connected faithful groups with each $\mathbb{C}^n$ is standard Lie theory. The substantive VERSF claim — and the load-bearing question a sceptic will ask — is: *why exactly these three channels, and why does this channel decomposition follow necessarily from fold geometry rather than being chosen to match a known answer?*

The programme argues (via axioms V1, GG2′–GG5) that the answer lies in the local degrees of freedom available at a 2D surface embedded in 3D: V1 asserts that the fold carries a single binary committed distinction; GG2′ asserts that local fold geometry must be characterised by independent degrees of freedom that are representationally complete at the boundary; GG3–GG5 progressively eliminate mixed or redundant decompositions. The three remaining independent channels are then the scalar distance to the embedding space (one real degree → $\mathbb{C}^1$), the orientation of the normal bundle (spinorial, two complex degrees → $\mathbb{C}^2$), and the extrinsic curvature of the embedding (→ $\mathbb{C}^3$).

Specific technical vulnerability: The assignment of the extrinsic curvature channel to $\mathbb{C}^3$ requires scrutiny. The second fundamental form of a 2D surface embedded in 3D is a symmetric 2 × 2 real tensor, which has three independent real components — not three complex ones. Getting to $\mathbb{C}^3$ therefore requires either (a) a complexification step that promotes the three real components to three complex degrees of freedom, or (b) a different characterisation of the curvature channel that yields complex dimension 3 directly. The programme must explain which of these is intended and why it is the right move rather than an ad hoc step to reach SU(3). This is the least obviously justified of the three channel assignments, and it yields the most non-trivial gauge group. A reader with differential geometry background will notice the gap immediately. Until the complexification or re-characterisation is made explicit and independently motivated, the $\mathbb{C}^3 \to$ SU(3) step should be treated as a claim, not a derivation.

The claim is that no other decomposition satisfies all of GG2′–GG5 simultaneously. This argument needs to be read in the full companion derivation to evaluate whether the axioms earn the decomposition or assume it; the summary here is a sketch, not a proof.

- *Structural result (conditional):* Given V1, GG2′–GG5, and the three-channel decomposition, SU(3) × SU(2) × U(1) is the minimal symmetry structure.
- *Conditional within the VERSF structural programme:* the V1/GG axiom set; the completeness and independence of the channel decomposition; and chirality ($SU(2)_l$), which remains open .

9.2 Democratic allocation and geometric coupling invariant

- *Structural result:* Under the "no extra bits" condition, curvature is democratically allocated among 12 generators, producing the geometric invariant α(geom) = $(1/12)^2$ = 1/144.
- *What the "no extra bits" argument consists of:* The argument runs as follows. The fold boundary has a fixed total geometric information content determined by its embedding constraints. Under the One-Fold conditions, the 12 generators of SU(3) × SU(2) × U(1) partition this content. "No extra bits" means no generator receives a preferred allocation — the total curvature is distributed equally. This is not a symmetry assumption about the physical coupling constants; it is a structural assertion about the information content of the fold boundary before any physical identification is made. The question a sceptic should ask is: *why is equal allocation the right condition rather than, say, allocation proportional to the dimension of each factor?* The programme does not currently give a uniqueness argument for equal allocation over other allocation rules at this level of detail, which is why the circularity risk below remains live.
- *What a convincing uniqueness argument would look like:* The programme needs a symmetry principle that treats all 12 generators equivalently at the fold-boundary level, *prior to and independent of* any physical identification of generators with coupling constants. Such a principle would need to show that no geometric property of the fold boundary distinguishes between generators before the identification step — effectively a fold-level permutation symmetry over the generator set. If such a symmetry can be derived from the fold axioms, equal allocation follows as the unique choice. If it cannot, the democratic allocation remains an additional assumption. Naming this gap makes the programme's next required step precise.

- *Critical caveat:* If the equal-allocation rule was selected specifically because it gives 1/144 ≈ α rather than because it follows from a prior principle, the derivation is circular. Asserting non-circularity is not the same as demonstrating it. The programme needs to show that equal allocation follows from a constraint that can be stated and evaluated independently of the numerical target.
- *Conditional within the VERSF structural programme:* democratic allocation axiom and One-Fold structural conditions.

9.3 Physical identification of α(geom) with the fine-structure constant

This result is currently a physical identification with three unfinished bridge steps, not a structural derivation. Readers should not weight the numerical agreement $\alpha^{-1} \approx 137.14$ as strong evidence for the framework until those bridges are complete.

- *Physical identification:* Mapping α(geom) = 1/144 to the physical electromagnetic coupling requires: (i) a gauge-coupling mapping from fold geometry to low-energy gauge theory; (ii) a $3 \oplus 1$ impedance correction; (iii) explicit RG running from the fold scale to low energies with no fitted intermediate parameters.
- *Status:* None of these three bridge steps has been derived within the programme. This is a definitive statement about the programme's current state, not a hedged one: the companion corpus has been surveyed and none of the three steps — gauge-coupling mapping, impedance correction, or RG running with no fitted parameters — appears in completed form. The programme's failure condition F6 (§13) explicitly recognises that if step (iii) requires a fitted parameter, the numerical agreement is not a prediction.

9.4 Cosmological constant scaling theorem Under the L1–L4 vacuum-energy axioms:

$$\Lambda/\Lambda(\text{Planck}) = C f^2$$

- *Structural result:* the scaling relation itself is derived from the axioms — this is the genuine achievement, giving a much tighter window than naïve QFT vacuum-energy estimates.
- *Physical identification:* Mapping Cf^2 Λ(Planck) to the observed Λ is a separate Class III identification. *Conditional within the VERSF structural programme — * L1–L4 axioms and the physical identification bridge.

---

10. Particle Ontology: Conditional Results

10.1 Four admissible physical fold states The fold carries one bit of committed distinguishability while admitting four admissible physical orientation states, because commitment polarity and boundary orientation are structurally independent.

*Conditional within the VERSF structural programme — * fold-boundary geometry.

10.2 Electrons as minimal stable closures

- *Structural result:* Electrons are minimal self-sustaining admissible fold-closures within the closure ontology.
- *Physical identification:* Mass as closure-stabilisation energy; spin and charge as emergent structural properties. Not yet precision-phenomenological.

10.3 Quarks as partial closures requiring composite stabilisation

- *Structural result:* Quarks are incomplete closures that do not close admissibly in isolation.
- *Physical identification:* This gives a structural confinement picture — isolated quarks are topologically inadmissible, not merely energetically suppressed.

10.4 Proton as three-channel composite closure

- *Structural result:* The proton is a stable three-channel composite fold-closure.
- *Physical identification:* Most of its mass arises from closure-stabilisation energy, consistent with the observed ratio of proton mass to current-quark masses.

---

11. TPB–BCB Comprehensive Programme

11.1 Tick as unit vortex on the void–universe interface

- *Structural result:* A tick is the creation or annihilation of a unit vortex on a hexagonally tiled void–universe interface. Appendix I of *Towards a Complete Information-Theoretic Physics* argues that vortices are the unique tick carriers satisfying locality, stability, discreteness, and isotropy.
- *Conditional within the VERSF structural programme — * the vortex microphysical axioms.

11.2 Landauer–CMB bit-energy matching

- *Physical identification:* The bit-energy scale is grounded in the Landauer bound at the CMB temperature, treating this as a thermodynamic boundary condition tying microphysics to cosmology.
- *Clarifying note:* This identification imports a specific cosmological assumption — that the CMB temperature is the relevant thermal environment at which the Landauer erasure cost should be evaluated. This is not a neutral or uniquely motivated choice: it assumes the universe's current large-scale thermal state sets the fundamental bit-energy scale, which is a non-trivial claim. The identification is treated as a boundary condition rather than a derived result within the programme, but readers should be aware that it anchors the microphysics to an empirically measured cosmological quantity, not to a quantity derived from the framework's axioms.

11.3 Role-4 field equations and GR recovery

- *Structural result:* The Extremal Distinguishability–Entropy Principle (EDEP) — that physical configurations extremise distinguishability gained per entropy produced — yields a complete set of coupled field equations for the entropy field s(x) and time-depth field τ(x), together with modified Einstein equations. A nontrivial Fisher-metric coefficient relation:

$$\xi_2^2 = \varepsilon^2 \xi_1 \kappa_4$$

is derived from 2×2 Fisher metric algebra, reducing free parameter count.

- *Physical identification:* In the low-gradient limit, Role-4 reduces exactly to General Relativity. Compatibility with the GW170817 gravitational-wave speed constraint is demonstrated.
- *Conditional within the VERSF structural programme — * EDEP and the Role-4 field structure.

11.4 Fermion Fold Principle and three stable generations

- *Structural result:* Fermion species correspond to topological minimisers of a Fisher-distinguishability functional on $\mathbb{CP}^2 \times \mathbb{CP}^1$. Exactly three stable fold configurations exist at winding numbers (1,0), (1,1), (2,1). The number three is treated as theorem-like within the framework.
- *Physical identification:* These three configurations are identified with the observed three fermion generations.
- *Conditional within the VERSF structural programme — * the FFP axioms; exact mass ratios require numerical solution .

11.5 Rigidity of the Yukawa sector Appendix L (*Towards a Complete Information-Theoretic Physics*) gives a master rigidity theorem: fold radii are unique, mass ordering is forced, and Yukawa structure is rigid in direction and spacing. Exact numerical masses await fold-equation solution.

*Conditional within the VERSF structural programme — * FFP framework. Numerical execution is open work .

---

12. Explicit Constructive Realisations

12.1 Standard Model from hexagonal geometry *(Class IV)* A companion construction provides an explicit microscopic realisation in which the conditional One-Fold structural conditions are dynamically satisfied. The emergent long-wavelength theory reproduces U(1), SU(2), SU(3), Higgs, and confinement sectors.

12.2 Binary fold substrate recovers the path integral *(Class IV)* The sum over binary fold configurations converges to the standard path integral for quadratic actions, with exact continuum-limit recovery for all free non-interacting quantum theories.

12.3 Tsirelson consistency in the shared-boundary entanglement model *(Class IV)* Bell-CHSH violation is derived from shared boundary geometry:

$$S_{max}(J) = 2\sqrt{1 + \sin^2 2J}$$

reaching the Tsirelson bound at $J = \pi/8$. Deviations are sub-Tsirelson, never super-Tsirelson. Local POVMs on disjoint fold domains commute, preserving no-signalling.

Part III — Failure Conditions and Falsification Tests

This tier identifies what remains undone and, crucially, what would constitute failure of the programme or specific sectors of it. A framework that cannot specify its own failure conditions is not a physical theory.

---

13. Programme Failure Conditions

The following are explicit conditions under which specific sectors or the programme as a whole would be considered to have failed. These are stated before the open-problem list so that the problem list can be read with failure stakes in mind.

F1. Interface law and fold boundary structure. The programme does not require a direct microphysics of the Void's interior — that is deliberately outside its epistemic scope. What it does require is a derivable interface law: explicit structural constraints governing the fold boundary at which reversible distinguishability becomes irreversible physical record. If no such interface law can be derived or specified that (a) generates commitment events consistently with Axioms A1–A3, (b) implies the admissibility filter, and (c) grounds the commitment barrier Φ_c, then the downstream programme fails. The failure condition is not "we cannot model the Void" but "we cannot characterise the fold."

F2. Commitment barrier derivation. If Φ_c cannot be derived from BCB/TPB substrate axioms, the rate law (§6.1) retains an unmotivated free parameter. If that parameter must be fitted to recover known physics, the fact-production equation loses its explanatory force and becomes a phenomenological fit.

F3. Complex inner-product structure. If the complex inner-product structure of commitment amplitudes cannot be derived from or independently justified by fold axioms, the deepest quantum route — including the Tsirelson derivation and gauge group classification — rests on an assumption rather than a theorem. This would not collapse those results, but it would prevent their elevation to Part I.

F4. Generic Lorentz universality class. If fold-admissible commitment structures can be shown to admit non-Lorentzian coarse-grained limits (other than as measure-zero special cases), the claim that Lorentz invariance is emergent rather than assumed would fail.

F5. Chirality. If the fold-boundary geometry cannot be shown to generate left-right asymmetry, the gauge group derivation remains stuck at SU(3) × SU(2) × U(1) rather than the chiral Standard Model. A non-chiral derivation of the Standard Model is not the Standard Model.

F6. Fine-structure constant derivation. If the RG bridge from α(geom) = 1/144 to the low-energy value cannot be completed without a fitted intermediate parameter, the numerical agreement $\alpha^{-1} \approx 137.14$ should be regarded as coincidence, not derivation.

F7. Fermion mass spectrum. If the FFP fold equations, when solved numerically, produce mass ratios inconsistent with observation outside the framework's stated theoretical uncertainties, the Fermion Fold Principle approach to the mass problem fails.

F8. Experimental exclusion. If direct laboratory tests (§14.1–14.3) are performed at the relevant precision and find no signal, those sectors are falsified. The strength of the framework's experimental programme is precisely that these are risky predictions, not post-hoc fits.

---

14. Experimental Programme

The experimental programme is divided into direct laboratory tests and indirect astrophysical or cosmological tests.

15A. Direct Laboratory Tests

15.1 Geometry-dependent decoherence modulation *(Highest priority)* The fold substrate model predicts that dephasing rates are modulated by fold angle α as $\tan^2\alpha$, that T_2 (transverse coherence time) is strain-dependent under fixed environmental coupling, and that decoherence rates vary with system geometry in ways not present in standard quantum mechanics.

The functional form $\tan^2\alpha$ gives a parametric prediction, but a numerical estimate of the magnitude of the modulation in experimentally accessible strained systems — and the strain and geometry ranges within which the effect would first exceed current measurement precision — has not yet been computed within the programme. Until that estimate is provided, the falsification condition below has the same limitation as §14.2: the prediction is in principle falsifiable but not yet practically so, because "the predicted sensitivity level" is not yet a determinate quantity.

- *Experimental target:* High-precision spin-coherence experiments in strained systems, comparing decoherence rates across geometric configurations while holding environmental coupling constant.
- *Falsification condition:* Absence of geometry-dependent modulation at the predicted sensitivity level would falsify this sector — *provided* the programme first specifies the predicted magnitude in experimentally accessible parameter ranges. Until then, this prediction is in principle falsifiable but requires the order-of-magnitude computation to become practically so.

15.2 Sub-Tsirelson reduction in distributed multipartite entanglement For multipartite entanglement distributed over large spatial domains, finite fold stiffness and energy budget impose a ceiling: maximum CHSH violation is reduced below $2\sqrt{2}$ as a function of separation scale and entanglement distribution. Standard quantum mechanics predicts no such reduction.

The programme derives the parametric form $S_{max}(J) = 2\sqrt{(1 + \sin^2 2J)}$, reaching $2\sqrt{2}$ only at $J = \pi/8$, with sub-Tsirelson values for other fold configurations. The effect is predicted to become detectable at spatial separations where the fold energy budget per entangled pair is comparable to the fold stiffness scale — which is related to the coherence scale $\xi \approx 8 \times 10^{-5}$ m. An order-of-magnitude estimate of the reduction magnitude and the separation threshold at which it first exceeds current experimental precision has not yet been computed within the programme; until it is, the falsification condition below has no practical precision threshold.

- *Experimental target:* Large-separation multipartite CHSH experiments with precision sufficient to detect sub-Tsirelson deviations. The relevant parameter regime is separations comparable to or exceeding the fold coherence scale, and entanglement distributions approaching the energy-budget limit.
- *Falsification condition:* Confirmed achievement of the full Tsirelson bound $S = 2\sqrt{2}$ at separations and entanglement levels within the predicted sensitive regime would falsify this prediction — *provided* the programme first specifies the predicted magnitude and threshold to a level that makes "within the sensitive regime" a determinate condition. Until the order-of-magnitude estimate is computed, this prediction is in principle falsifiable but not yet practically so.

15.3 Coherence-scale effects in Casimir-type experiments The quartic capacity relation predicts a characteristic coherence scale $\xi \approx 8 \times 10^{-5}$ m as the minimum record-forming region. Casimir-type experiments probing vacuum fluctuations at or near this scale should exhibit deviations from standard QFT predictions.

- *Experimental target:* Precision Casimir measurements in the 10^{-4} m regime, or experiments designed to probe vacuum structure at the predicted coherence length.
- *Falsification condition:* Absence of any anomalous signal in experiments reaching the relevant sensitivity at this scale would challenge the physical identification of ξ.

15B. Indirect Astrophysical and Cosmological Tests

15.4 CMB small-scale structure and early-galaxy formation statistics If the distinguishability capacity χ(L) is physically real, finite capacity should produce observable deviations in the small-scale CMB power spectrum and in early-galaxy formation statistics relative to standard inflationary predictions.

- *Experimental target:* Next-generation CMB surveys (CMB-S4 and successors) and high-redshift galaxy surveys capable of resolving the predicted structure.
- *Falsification condition:* Precise agreement with standard inflationary predictions at all accessible scales, with no residual signal at the coherence scale, would constitute pressure on the physical identification of ξ.

15.5 Gravitational wave propagation in Role-4 gradient regimes Role-4 field equations reduce to GR in the low-gradient limit (consistent with GW170817, $|v(GW) - c|/c \lesssim 10^{-15}$). In regimes where Role-4 and GR differ — strong entropy-field gradients, near structured matter distributions — gravitational wave speed or waveform should deviate.

- *Experimental target:* Third-generation gravitational wave detectors (Einstein Telescope, Cosmic Explorer) with sensitivity to waveform deviations in strong-gradient environments.
- *Falsification condition:* Precise agreement with GR predictions in those gradient regimes where Role-4 predicts measurable deviation would falsify the Role-4 field structure.

15.6 Fermion mass spectrum as internal testability criterion Once the FFP fold equations are solved numerically, the predicted mass spectrum either agrees or disagrees with measured quark and lepton masses. This is an internal testability criterion requiring no new experiment — only completion of the numerical programme.

- *Falsification condition:* Predicted mass ratios outside observational constraints by more than the framework's stated theoretical uncertainties would falsify the Fermion Fold Principle approach.

Appendix: Programme-Level Structural Achievements

Beyond individual results, the VERSF programme has achieved several structural advances at the level of how the research is organised.

Non-emptiness of constraint sets. The programme demonstrates that its structural conditions are satisfiable: at least one explicit microscopic or mesoscopic construction (hexagonal geometry, binary fold substrate, vortex tick model) satisfies the programme's constraints. This prevents the framework from being a purely verbal theory with no realisation.

Three-tier audited structure. *Old-Fold-Companion 5* explicitly separates unconditional theorems, conditional structural derivations, and physical identifications. This level of internal self-criticism is unusual in speculative theoretical physics and is one reason the programme can be assessed rigorously.

Explicit failure conditions and skeptic's checklist. *Towards a Complete Information-Theoretic Physics* ends with a detailed response to objections — why the Tsirelson bound, why three generations, why these masses, why this is not modified-gravity crankery — with each objection directed to a specific derivation. The programme failure conditions in §13 of this document extend that culture to the full programme.

Cross-paper internal consistency. The architecture paper explicitly aligns notation, defines three theory layers, fixes vocabulary, and works to prevent layer-mixing across a large and growing corpus. Maintaining internal consistency across the programme's more than one hundred companion papers is a genuine organisational achievement that makes external assessment like this document possible.