Structural Necessity of Lorentzian Geometry in VERSF

Sequential Commitment Transport, Refinement Stability, and the Emergence of Covariant Geometric Structure

Keith Taylor — VERSF Theoretical Physics Programme

General-Reader Summary

Modern physics treats Lorentzian geometry as a starting point — a 4-manifold with a metric of signature (+, −, −, −), a universal maximum speed, observer-dependent time, and causal organisation by light cones. Einstein's relativity takes these as postulates and works downstream from them.

This paper asks the upstream question: *why should reality organise itself this way at all?*

In VERSF, neither space nor time is fundamental. The universe is built from irreversible commitment events — distinguishable records forming on a finite substrate. Space, time, and the geometry that binds them together appear only after coarse-graining large collections of such events.

The central claim is that Lorentzian geometry is the **unique stable large-scale geometry** compatible with five substrate-level conditions:

1. finite propagation,
2. observer-invariant distinguishability (A0),
3. operational inaccessibility of any substrate rest frame,
4. refinement-compatible sequential transport,
5. homogeneity and isotropy of admissible coarse-graining.

The logical core is short. The cone field on the coarse-grained manifold is forced by finite propagation plus operational substrate-frame inaccessibility. By a classical result of Malament and Hawking–King–McCarthy, the cone field then determines the metric *up to a conformal factor*. The conformal factor is fixed (up to overall units) by transport-density preservation, which is itself a consequence of substrate-level bit conservation (BCB). At each point of the coarse-grained manifold, the structure group is therefore the proper orthochronous Lorentz group $SO^{+}(1,3)$. When the substrate produces a *flat* continuum limit, the global symmetry group is the proper orthochronous Poincaré group $\mathcal{P}^{+}$ — and the Alexandrov–Zeeman theorem then identifies it as the full causal-automorphism group of Minkowski space.

Minkowski geometry is therefore not primitive. It is the only stable large-scale bookkeeping structure compatible with finite distinguishability and invariant transport. Relativity ceases to be a postulate about geometry; it becomes a consistency condition on how irreversible facts can propagate through a finite substrate.

This paper also closes a major dependency in the broader BCB–VERSF programme. Earlier papers derived Hilbert-space structure, gauge connections, SU(3) × SU(2) × U(1) gauge structure, chirality, and hypercharge constraints from distinguishability conservation and Fisher geometry. However, those derivations still relied implicitly on a Lorentzian continuum and spinorial covariance. The present paper supplies that missing geometric layer by deriving local Lorentz structure itself from finite distinguishability, transport-cone invariance, and substrate-frame gauge redundancy. In this way, the Standard Model reconstruction programme and the emergent-geometry programme become part of a single constrained architecture rather than parallel derivation tracks.

Abstract

We derive Lorentzian geometric structure within the Void Energy–Regulated Space Framework (VERSF) from five substrate-level inputs: finite propagation, observer-invariant distinguishability (A0), irreversible commitment dynamics (A2), refinement-compatible sequential transport, and admissible coarse-graining. Earlier VERSF papers established a finite substrate propagation speed, observer-protocol invariance, persistent cohomological transport, Maxwell-form U(1) dynamics, spinorial closure, CAR-algebra emergence, and relativistic fermionic field structure. The unresolved structural question was whether Lorentzian geometry itself is emergent or merely inherited.

This paper also provides the missing geometric-covariance layer required by the BCB derivation of gauge structure and Standard Model internal symmetry. Earlier BCB papers derived gauge connections, SU(3) × SU(2) × U(1) structure, chirality constraints, and hypercharge uniqueness from distinguishability conservation, Fisher geometry, anomaly exclusion, and entropy minimization, but still relied implicitly on a Lorentzian continuum. The present paper closes that dependency by deriving local Lorentzian geometry from substrate transport consistency itself.

We prove four results, each labelled with its epistemic status.

- **Theorem 1L (Local Lorentz Emergence — proven, conditional on H1–H5 and Proposition T1).** At each point x of the coarse-grained manifold $\mathscr{M}$_coarse, the structure group of T_x $\mathscr{M}$_coarse compatible with the cone, transport-density, orientation, and time-orientation conditions has connected component $SO^+(1,3)$; the full admissible structure group is at least $O\uparrow(1,3) = SO^+(1,3) \rtimes \{1, P\}$.
- **Theorem 1G (Global Poincaré Emergence — proven, conditional on H1–H5 plus flat continuum limit).** When the substrate produces a flat continuum limit — i.e., when $\mathscr{M}$_coarse admits a flat affine connection compatible with the local $SO^+(1,3)$ structure —

the connected component of the global admissible-transformation group is the proper orthochronous Poincaré group $\mathcal{P}^{+} = SO^{+}(1,3) \ltimes \mathbb{R}^{1,3}$.

- **Theorem 2 (Invariant Interval — proven, conditional on the hypotheses of 1L).** The unique (up to overall scale) quadratic form on T_x $\mathcal{M}$_coarse invariant under $SO^{+}(1,3)$ is

  $$ds^2 = c^2\, dt^2 - dx^2 - dy^2 - dz^2.$$

- **Theorem 3 (Boost Generation from Refinement — proven, conditional on H1–H5, T1, and verification of R1).** Lorentz boosts arise as the one-parameter families of admissible transformations preserving the cone, transport density, and refinement limits of the sequential-transport functor. The boost parameter coincides with the rapidity $\eta = \operatorname{arctanh}(v/c)$.

A key methodological move — and the principal sharpening over the v1 draft — is the explicit separation of three steps that were previously fused. *(a)* Malament (1977) and Hawking–King–McCarthy (1976) supply the cone-to-conformal-class step on the strongly-causal time-orientable Lorentzian continuum produced by admissible coarse-graining (Lemma 5.2). *(b)* Transport-density preservation (now a derived proposition, not an axiom) fixes the conformal factor. *(c)* Alexandrov (1950s) and Zeeman (1964) characterise the global symmetry group on the flat-Minkowski continuum limit. We also address the potential circularity in the operational-undetectability argument by giving a *purely* operational characterisation of admissibility, with no residual appeal to A0.

The result converts Lorentz covariance from an inherited continuum assumption into a structural consequence of finite distinguishability and commitment transport. The transport-cone structure also supplies the missing geometric layer underlying the emergent Dirac-field programme. Microcausal suppression follows naturally at the geometric level, though full *algebraic* microcausality requires one further structural ingredient identified and now precisely stated in §11. We close by identifying the remaining open problems — exact continuum-limit control, substrate-level derivation of H3 and H4, interacting gauge-field emergence, renormalisation closure, and full Standard Model reconstruction.

Notation glossary

Notation is as in the sequential-interface-transport, σ-duality, BCB–VERSF synthesis, and admissible-coarse-graining papers, with the following recap:

Symbol	Meaning
$\mathcal{M}$_coarse	4-dim connected smooth coarse-grained manifold
$\mathcal{S}$_substrate	Substrate state space (pre-coarse-graining)
$\mathcal{C}(x)$	Causal transport cone at $x \in \mathcal{M}$_coarse
$\mathcal{C}^{+}(x)$, $\mathcal{C}^{-}(x)$	Future and past lobes of $\mathcal{C}(x)$

Symbol	Meaning
ρ_T	Coarse-grained transport-current density, $\rho_T = J^\mu n_\mu$
J^μ	Coarse-grained transport current (substrate-level bit-current)
$\mathcal{R}$: $TPB_k \to TPB_{k+1}$	Refinement functor on Ticks-Per-Bit layers
$\mathcal{O}_{adm}$	Space of admissible observer-protocols (defined §4.2)
$\mathcal{P}^+$	Proper orthochronous Poincaré group $SO^+(1,3) \ltimes \mathbb{R}^{1,3}$
$O{\uparrow}(1,3)$	Orthochronous Lorentz group $SO^+(1,3) \rtimes \{1, P\}$
BCB	Bit Conservation and Balance
TPB	Ticks-Per-Bit

Forward reference: the substrate-to-geometry dependency diagram in §12 may be useful to consult before reading the main results.

Table of Contents

---

1. Introduction

Special relativity is usually introduced axiomatically — inertial-frame equivalence, invariance of c, Lorentzian geometry — with all consequences derived downstream. VERSF inverts the order. It begins from finite distinguishability, irreversible commitment, local transport, and observer-invariant content, and asks what large-scale geometry must hold if those substrate principles do.

Previous VERSF papers established information-theoretic, Hilbert-space, continuous-reversible-evolution, spinorial, fermionic, and relativistic-field structure from substrate-level principles. They also established that the substrate admits a finite invariant propagation speed and Lorentz-compatible transport cones. What they did *not* establish — and what this paper addresses — is the converse direction: *why Lorentzian geometry is the only stable choice* given the substrate constraints.

The importance of this gap increased substantially after the development of the BCB gauge-structure programme. In those papers, gauge connections were shown to arise necessarily from local state comparison and distinguishability conservation on Fisher-information manifolds, while SU(3) × SU(2) × U(1) structure emerged from projective probability geometry, entropy bounds, anomaly constraints, and representation-theoretic minimality. However, those constructions still presupposed a Lorentz-compatible continuum framework in which spinors, gauge currents, and local field transport could be consistently defined. The present paper therefore serves as the geometric closure layer linking the BCB gauge programme to the emergent-geometry programme.

The central question is:

Why must the large-scale geometry of admissible transport be specifically Lorentzian, rather than Galilean, Euclidean, Carrollian, Finsler, doubly-special-relativistic, or Hořava–Lifshitz?

The answer developed here is that finite propagation, together with the impossibility of operationally detecting substrate-frame motion, forces invariant causal transport cones; that the metric is then forced (up to overall units) by transport-density preservation; and that boost structure is forced by refinement compatibility. Relativity emerges not as a geometric axiom but as a transport-consistency theorem.

We note in §14 that the residual mathematical risk in this derivation lies not in the Lorentz-emergence step itself — which, given a smooth non-degenerate cone field and the bit-conservation properties of BCB, is mechanically tight — but in the continuum-limit regularity step that supplies the smooth cone field assumed throughout. The structural emergence of Lorentzian geometry from substrate principles is therefore on firmer ground than the prerequisite continuum-limit theorem; readers tracking the open problems should keep this asymmetry in mind.

---

2. Hypotheses, axioms, and propositions

To make the dependency structure explicit we list the inputs separately from the substrate axioms A0–A4 of the VERSF programme.

2.1 Substrate axioms (carried in from prior VERSF papers)

- **A0 — Observer-invariant distinguishability.** Physical content depends only on distinctions invariant under all admissible observer-protocols.
- **A1 — Finite distinguishability.** Any bounded substrate region admits only finitely many operationally distinguishable states.
- **A2 — Irreversible commitment.** Facts correspond to irreversible commitment events.
- **A3 — Reversible pre-commitment evolution.** Between commitment events, admissible transport evolution is reversible.
- **A4 — Local coupling.** Substrate interactions propagate only between neighbouring admissible states.
- **BCB — Bit Conservation and Balance** *(carried in from the BCB–VERSF synthesis paper).* Substrate bit content is conserved and balanced along admissible transport sequences; equivalently, the substrate bit-current is divergence-free at the coarse-grained level.

2.2 Hypotheses used in this paper

- **H1 — Finite invariant propagation.** There exists a finite $c > 0$ such that no admissible commitment-transport sequence exceeds spatial rate c in the coarse-grained limit. *(Theorem of prior VERSF work from A1 and A4; carried in here as input.)*
- **H2 — Operational inaccessibility of substrate frame.** No admissible observer-protocol $O \in \mathcal{O}\_adm$ produces an output dependent on the substrate's update foliation $\mathcal{F}_0$ when restricted to invariant content. *(§4.2 gives the purely operational characterisation of admissibility that prevents circularity.)*
- **H3 — Homogeneity of admissible coarse-graining.** The space of admissible coarse-graining maps φ: $\mathcal{S}$_substrate → $\mathcal{M}$_coarse is invariant under spatial translations on $\mathcal{M}$_coarse.
- **H4 — Isotropy of admissible coarse-graining.** The space of admissible coarse-graining maps is invariant under rotations on $\mathcal{M}$_coarse.
- **H5 — Refinement compatibility.** Sequential transport commutes with refinement of the TPB layer (precise commutation diagram given in §8.1).

2.3 Propositions derived in this paper (formerly axioms)

- **Proposition T1 — Transport-density preservation.** *Admissible observer transformations preserve the coarse-grained transport-current density ρ_T = J^μ n_μ up to a global constant of units.* Derived in §6.1 from BCB and A0. This replaces the v1 axiom T1.
- **Proposition R1 — Refinement-stability of boosts.** *The one-parameter subgroups of G consisting of cone- and transport-density-preserving transformations that are not pure rotations or translations intertwine with the refinement functor $\mathcal{R}$.* This is either a self-consistency theorem (if the substrate-level proof in the sequential-interface-transport paper is invoked) or a labelled conjecture pending full substrate-level closure. We treat it explicitly as the former in §8 with the dependency flagged.

2.4 Epistemic status

A0–A4 and BCB are foundational VERSF axioms. H1 is a theorem of prior VERSF work. H2 is the load-bearing hypothesis of this paper — §4 is devoted to justifying it without circularity and §13 quantifies its empirical bounds. H3 and H4 are conditional assumptions on the coarse-graining procedure, flagged for future substrate-level derivation. H5 is a structural assumption on the refinement functor. T1 and R1 are now derived results, not axioms.

2.5 Relationship to the BCB gauge programme

The present paper should be read as structurally complementary to the BCB gauge-derivation papers. Throughout this paper, "BCB gauge programme" and "BCB gauge-derivation papers" refer to the relevant subsections of the BCB–VERSF synthesis paper (Hilbert reconstruction; local state-comparison and gauge-connection emergence; U(1) phase-redundancy derivation; SU(3) uniqueness and 3-body singlets; SU(2) chirality and fold-orientation; hypercharge structure) rather than to separate standalone papers. Where a result requires standalone treatment, it is named explicitly.

Those derivations established:

- Hilbert-space reconstruction from distinguishability preservation,
- projective state-space geometry,
- SU(n) emergence from volume-preserving isometries,
- Fisher-information geometry on $\mathbb{CP}^{n-1}$,
- gauge-connection necessity from local state comparison,
- U(1) emergence from phase redundancy,
- SU(3) uniqueness from 3-body singlets and entropy-capacity bounds,
- SU(2) chirality structure from fold orientation and spinor geometry,
- partial reconstruction of Standard Model hypercharge structure.

However, those derivations relied on a Lorentz-compatible continuum background supporting spinor representations, local field transport, gauge-covariant derivatives, causal propagation, and relativistic covariance.

The present paper supplies that missing geometric layer. In particular:

- Theorem 1L provides the local Lorentz structure required by spinorial and gauge-covariant constructions.
- Theorem 2 supplies the invariant interval underlying relativistic field propagation.
- Theorem 3 supplies the boost structure required for observer-independent transport covariance.
- The transport-cone structure provides the causal foundation underlying the local-comparison and gauge-connection arguments of the BCB framework.

Thus the BCB gauge programme and the emergent-geometry programme are no longer independent derivation tracks but become mutually constraining components of a single substrate architecture.

---

3. Finite propagation and sequential commitment transport

By A1 and A4, transport propagation is finite. Let ℓ denote the maximal substrate propagation distance per update step and ν the maximal substrate update rate. Then

$c = \ell \cdot \nu$

defines the maximal admissible transport speed. This speed is not a geometric postulate but a substrate throughput constraint — a ceiling on how fast distinguishable records can propagate before commitment dynamics ceases to be consistent with A2.

Sequential commitment transport (developed in the σ-duality and sequential-interface-transport papers) thereby defines a causal accessibility structure on the set of commitment events, a maximal propagation cone $\mathcal{C} \subset \mathscr{M}\_\text{coarse} \times \mathscr{M}\_\text{coarse}$, and a finite coordination rate for distinguishable updates.

We write $x \prec y$ when there exists an admissible commitment-transport sequence taking x to y, and $x \ll y$ for the strict relation. The pair $(\mathscr{M}\_\text{coarse}, \prec)$ is the *causal pre-order* induced by transport. Throughout, $\mathscr{M}\_\text{coarse}$ is taken to be a 4-dimensional connected smooth manifold; dimensional emergence is treated in the dimensional-emergence paper and is not re-derived here.

---

4. Observer-invariant distinguishability and substrate-frame gauge

This section is the load-bearing argument of the paper. We must justify H2 without circularity. The v2 fix relative to v1 is the elimination of any A0-reference from the operational characterisation of admissibility.

4.1 Why a naive statement is circular

A naive version of H2 says: *"no admissible observer detects the substrate frame"*. But if "admissible" is defined as "respecting Lorentz transport-cone structure", the argument collapses — we have built the conclusion into the definition of admissibility.

4.2 Purely operational characterisation of admissibility

An observer-protocol O is admissible iff:

1. **(Locality)** O is implementable as a finite composition of local substrate operations satisfying A4.

2. **(Commitment respect)** O produces outputs only through irreversible commitment events (A2).
3. **(Finite distinguishability)** O distinguishes only finitely many outcomes in any bounded region (A1).
4. **(Inter-protocol reproducibility)** The outputs of O are reproducible by any other protocol O′ satisfying clauses 1–3, modulo permutations of operationally indistinguishable outcomes.

Clause 4 replaces v1's "A0-respect" clause and makes no reference to A0. It is a closure condition on the protocol class: an admissible protocol is one whose outputs are not protocol-specific artefacts but instead lie in the common image of all protocols satisfying clauses 1–3.

4.3 Recovery of A0

A0 is then a *theorem* about this protocol class: the invariant content extracted by $\mathcal{O}$_adm coincides with the distinctions invariant under all admissible observer-protocols, by construction of clause 4. This is the genuinely circular-free version.

4.4 The forcing argument

Suppose a substrate foliation $\mathcal{F}_0$ exists. Either:

- *(i)* some admissible O can produce an output dependent on $\mathcal{F}_0$; or
- *(ii)* no admissible O can.

If *(i)* holds, then $\mathcal{F}_0$ is operationally detectable. This is an empirically falsifiable claim — current bounds on preferred-frame parameters from atomic-clock comparisons and astrophysical Lorentz tests sit at the 10^{-18}–10^{-20} level (or tighter in current editions; see §13).

If *(ii)* holds, then by clause 4 and the recovered A0, the distinction "moving relative to $\mathcal{F}_0$" vs "at rest relative to $\mathcal{F}_0$" is gauge — it carries no invariant content. The substrate may have a preferred update ordering, but no admissible observer can interact with it.

H2 is the statement that *(ii)* holds. It is *conditional* on the empirical bounds in §13 but no longer *circular*. H2 is therefore defeasible — any future preferred-frame detection above the §13 thresholds would falsify the hypothesis and, with it, the entire geometry-emergence chain (Lemma 5.1 onward). The information-geometric chain (Layer 3 of the §12 diagram: Fisher geometry, Hilbert structure, projective probability geometry) does not depend on H2 and would survive an F1 falsification; what would collapse is the geometric layer downstream of Lemma 5.1.

---

5. Emergent causal cones

Finite propagation defines, at each point x ∈ 𝓜_coarse, the set

$\mathcal{C}$(x) = { y ∈ 𝓜_coarse : x ≺ y or y ≺ x }

of events transport-connected to x. The complement is the set of transport-incompatible events.

Lemma 5.1 (Cone invariance). *Under H2, every admissible observer agrees on the partition ($\mathcal{C}$(x), 𝓜_coarse ∖ $\mathcal{C}$(x)) at every x.*

Proof. If two admissible observers O, O′ disagreed on whether y ∈ $\mathcal{C}$(x), the partition itself would depend on the observer. But the partition is defined operationally by the existence of an admissible commitment-transport sequence connecting x and y, and by clause 4 of the §4.2 admissibility characterisation, this existence is observer-invariant — any admissible protocol must reproduce the same accessibility verdict.

This is the substrate origin of light-cone structure. Lemma 5.1 is the proper statement of what was asserted (without proof) in earlier drafts as "causal cones are invariant."

We henceforth assume $\mathcal{C}$(x) is a non-degenerate open convex double cone at every x. The regularity properties needed are stated and proved as Lemma 5.2.

Lemma 5.2 — Non-degenerate cone regularity. *Under H1, H2, H3, and H4, the invariant transport cone at each point is a non-degenerate, open, convex, double cone.*

Proof. H1 excludes the two degenerate limits: $c = \infty$, which would remove a finite cone and produce Galilean structure, and $c = 0$, which would collapse propagation to Carrollian structure. H2 makes the accessibility partition observer-invariant; equivalently, the cone field is well-defined on 𝓜_coarse rather than only on individual observer slicings. H3 and H4 remove preferred spatial locations and directions from the admissible coarse-grained cone, so the cone at x is invariant under spatial translations and rotations. Convexity is a continuum-regularity property: the substantive content is that the continuum limit admits closure under infinitesimal composition of admissible transport directions. Under the smooth coarse-graining assumption, this induces convexity of the tangent cone. The required continuous interpolation in the continuum limit is itself a continuum-regularity assumption — formally established (where it holds) by the convexity-of-admissible-directions theorem of the sequential-interface-transport paper, and flagged in §14 as part of the broader continuum-limit regularity programme. A2 (irreversibility) selects the future lobe globally. Subject to that regularity, the cone is therefore open, convex, non-degenerate, and double-lobed. This completes the proof of Lemma 5.2.

Lemma 5.2 supplies the regularity conditions assumed by the Malament–Hawking–King–McCarthy theorem invoked in §6.2: the global pair (𝓜_coarse, ≺) is *strongly causal* in the standard sense, and the cone field is sufficiently regular for the cone-to-conformal-class result to apply.

5.3 Scope of the Malament–HKM invocation

Clarification on logical scope. The present paper does not invoke Malament–Hawking–King–McCarthy as an independent derivation of Lorentzian geometry from arbitrary substrate dynamics. Rather, Lemma 5.2 establishes that admissible coarse-grained transport defines a strongly causal cone structure on a smooth continuum limit. Malament–HKM is then applied conditionally: among continuum geometries admitting such a cone structure, the causal order uniquely fixes the conformal metric class.

The logical structure is therefore:

1. substrate transport + admissibility → invariant non-degenerate cone structure;
2. cone regularity + smooth continuum limit → conformal pseudo-Riemannian metric class with the cone as its null cone (Malament–HKM); together with the (1,3) signature that follows from Lemma 5.2 applied to the 4-dimensional $\mathscr{M}$_coarse of §3, and the Finsler-exclusion argument of §6.10, the class is conformal Lorentzian;
3. BCB transport-volume preservation → fixing of the conformal factor.

The claim of the paper is therefore not that Lorentzian geometry appears *ex nihilo*, but that once a smooth causal continuum limit exists, Lorentzian structure is the unique stable admissible geometry compatible with the substrate constraints. The role of Malament–HKM is therefore classificatory rather than generative: it classifies which continuum geometries are compatible with a given strongly-causal cone structure, rather than generating a Lorentzian manifold from non-Lorentzian inputs.

---

6. Theorems 1L and 1G — Lorentz and Poincaré emergence

We now prove the main symmetry result. The v2 sharpening separates the *local* (tangent-space) and *global* (manifold-wide) statements, and corrects the v1 misattribution of Alexandrov–Zeeman.

6.1 Proposition T1 — Transport-density preservation from BCB

Before stating Theorem 1, we derive what was previously axiomatised. The §6.1 derivation chain runs: BCB → divergence-free current (Lemma 6.1.a) → integrated charge invariance up to scale function λ (Lemma 6.1.b) → λ is a global constant (Lemma 6.1.c) → Proposition T1, which is the conjunction of 6.1.a–c (transport-density preservation) → transport-volume preservation, local form (Lemma 6.1.d) → conformal factor fixed (Corollary 6.1.e). §6.2 then uses Corollary 6.1.e to upgrade the conformal Lorentzian class to a metric Lorentzian structure in Step C of the Theorem 1L proof.

Lemma 6.1.a — Coarse-grained current conservation. *Under BCB, the substrate-level bit-current admits a coarse-grained continuum limit J^μ on $\mathscr{M}$_coarse satisfying ∂_μ J^μ = 0.*

Proof. BCB asserts that substrate bit content is conserved and balanced along admissible transport sequences (A4-local couplings). Under the admissible coarse-graining map φ:

$\mathcal{S}$_substrate → $\mathcal{M}$_coarse (developed in the admissible-coarse-graining paper), the discrete substrate bit-current pushes forward to a continuum current J^μ on $T_x \mathcal{M}$_coarse. Conservation lifts to divergence-freeness $\partial_\mu J^\mu = 0$ by the standard substrate-to-continuum limit theorem of that paper. This completes the proof of Lemma 6.1.a.

Lemma 6.1.b — Integrated bit count is admissible-observer-invariant. *Under A0, the §4.2 admissibility characterisation, and Lemma 6.1.a, any admissible observer transformation preserves the integrated bit count $Q(\Sigma) = \int_\Sigma J^\mu n_\mu \, d\Sigma$ over any admissible Cauchy surface Σ, up to a possibly position-dependent global scale function $\lambda: \mathcal{M}\text{\_coarse} \to \mathbb{R}^+$ relating bit-counting conventions between observers.*

Proof. By Lemma 6.1.a, Q(Σ) is independent of the choice of Cauchy surface Σ — this is the standard topological-invariance argument for conserved-current charges. By clause 4 of §4.2 (inter-protocol reproducibility), Q is invariant content: any two admissible protocols must agree on Q up to their respective bit-counting unit conventions. The relationship between two observers' bit-counts is, a priori, given by a positive scale function λ(x) — different regions of $\mathcal{M}$_coarse could in principle correspond to different unit conventions. The lemma asserts only that Q is preserved up to such a λ; it does not yet restrict λ to a constant. This completes the proof of Lemma 6.1.b.

Lemma 6.1.c — Finite distinguishability fixes the scale. *Under A1 (finite distinguishability), the scale function λ in Lemma 6.1.b reduces to a global constant.*

Proof. Suppose for contradiction λ is non-constant. Then there exist regions $R_1, R_2 \subset \mathcal{M}$_coarse with $\lambda(R_1) \neq \lambda(R_2)$. An admissible observer comparing bit counts across R_1 and R_2 would obtain operationally distinguishable rescalings — bit counts in R_1 would translate into bit counts in R_2 with a different conversion factor than the converse. By the locality clause (clause 1 of §4.2), this comparison is implementable by a finite composition of local substrate operations. By A1, the total bit count in any bounded region is finite and absolute. A position-dependent λ would therefore manufacture distinguishable bit-count assignments in a bounded region without corresponding substrate-level distinguishability, contradicting A1. Hence λ is constant. The constant is the choice of bit-counting unit, conventionally absorbed. This completes the proof of Lemma 6.1.c.

Proposition T1 (Transport-density preservation). *Under BCB, A0, and A1, admissible observer transformations preserve the coarse-grained transport-current density $\rho_T = J^\mu n_\mu$ up to a global constant of units.*

Proof. Combining Lemmas 6.1.a–c: the bit-current is divergence-free (6.1.a); the integrated bit count is admissible-observer-invariant up to a scale function λ (6.1.b); and λ is a global constant (6.1.c). The differential statement that $\rho_T = J^\mu n_\mu$ is preserved up to a global constant follows from the integrated statement by differentiating under the integral, using the smoothness of admissible observer transformations on the coarse-grained limit.

Lemma 6.1.d — Transport-volume preservation (local form). *Under BCB, A1, A0, and Proposition T1, for every $x \in \mathcal{M}$_coarse there exists an admissible normal neighbourhood U_x*

⊂ $\mathcal{M}$_coarse in which admissible observer transformations preserve the transport four-volume element associated with the coarse-grained current, up to one local constant.

Proof. BCB gives a conserved substrate bit-current whose coarse-grained limit satisfies $\partial_\mu J^\mu = 0$ (Lemma 6.1.a). Proposition T1 establishes that the transport density $\rho_T = J^\mu n_\mu$ is preserved up to a single global unit constant. In any admissible normal neighbourhood U_x, transport tubes $\mathcal{U} \subset U_x$ can be constructed locally as the union of integral curves of n_μ over a small Cauchy section: such tubes have zero lateral transport flux by construction, because in U_x the foliation by integral curves of n_μ is unobstructed. For any such tube bounded by admissible Cauchy sections Σ_1 and Σ_2,

$$Q(\mathcal{U}) = \int_\Sigma J^\mu n_\mu \, d\Sigma$$

is independent of the section Σ. Since $Q(\mathcal{U})$ is an invariant count of committed substrate distinctions, A1 forbids any local observer-dependent rescaling of the volume element within U_x: such a rescaling would change the number of distinguishable committed records assigned to subregions of $\mathcal{U}$ without changing the underlying substrate content. Therefore the measure $\rho_T \, d\Sigma$, and equivalently the induced transport four-volume element, is preserved on U_x up to a local constant. This completes the proof of Lemma 6.1.d.

Remark on globality. The lateral-no-flux tube construction is unconditional locally but is *not* in general unconditional globally — in curved continuum limits, the global existence of a sufficiently rich foliation by no-flux tubes is a regularity assumption about J^μ, formally established (where it holds) by the foliation lemma of the admissible-coarse-graining paper. Lemma 6.1.d is therefore stated and used in its local form throughout this paper; global constancy of the conformal factor in Corollary 6.1.e then follows from local constancy via connectedness of $\mathcal{M}$_coarse and continuity of Ω.

Corollary 6.1.e — Conformal factor fixed. *The conformal freedom $g_{\mu\nu} \to \Omega^2(x)\, g_{\mu\nu}$ left by causal-cone structure is reduced to Ω = constant by transport-volume preservation: locally by Lemma 6.1.d, and globally by connectedness and continuity of Ω on $\mathcal{M}$_coarse.*

Proof. In four dimensions, a conformal transformation rescales the metric volume element by $\Omega^4(x)$, and rescales hypersurface measures accordingly. By Lemma 6.1.d, in any admissible normal neighbourhood U_x the same conserved transport tube must be assigned the same local committed-record density by any two admissible observers. If $\Omega(x)$ varied within U_x, the integrand $\rho_T \, d\Sigma$ would transform under $\Omega^4(x)$ while the integral $Q(\mathcal{U})$ is fixed, forcing local density discrepancies. That contradicts Lemma 6.1.d, so Ω is constant on every U_x. Since $\mathcal{M}$_coarse is connected and Ω is continuous, local constancy on every U_x implies global constancy. The remaining constant is not physical geometry but a choice of units, absorbed into c and the normalisation of the metric. The causal conformal class [g] is therefore promoted to a unique Lorentzian metric up to global scale.

Corollary 6.1.e supplies the missing bridge between "cone determines conformal metric" (Malament–HKM, §6.2 Step A) and "transport fixes the metric" (the substrate-level reason the conformal factor is constant). This is the key proof upgrade: cone structure alone gives only a conformal class; the metric is fixed by the additional VERSF claim that committed distinctions are conserved substrate records, not observer-dependent bookkeeping artefacts. BCB together with A1 therefore does more than conserve information — through Lemma 6.1.d and Corollary 6.1.e, the joint content of substrate bit-conservation and finite distinguishability fixes the conformal scale of the emergent geometry.

Interpretive remark — relation to the BCB gauge programme. The transport-density current J^μ derived above admits a second interpretation within the broader BCB programme. In the BCB gauge-derivation work (located as subsections of the BCB–VERSF synthesis paper, including the Hilbert reconstruction, local-comparison, U(1) emergence, and SU(3) uniqueness arguments), distinguishability conservation on Fisher manifolds is shown to force gauge-covariant transport and principal-bundle connection structure.

Both derivations invoke substrate-level distinguishability currents defined at the TPB layer. Their commensurability — i.e., that the bit-current J^μ of this paper and the Fisher-information current of the BCB gauge-derivation are two coarse-grainings of a common substrate quantity rather than mathematically distinct objects — is the natural conjecture, but is not formally established in the present paper and is deferred to a dedicated identification result.

This is *not* a circularity even under the weaker (un-identified) reading. Both derivations invoke their respective substrate-level distinguishability currents at the TPB layer, prior to any continuum-geometric structure. The present paper derives Lorentzian structure from the transport properties of the bit-current; the BCB gauge programme derives gauge structure from the information-geometric properties of the Fisher-current. Neither derivation presupposes the other's continuum output. The fact that both substrate-level currents support downstream structures in the same programme is the architectural content that makes the programme convergent rather than merely concurrent — and identifying the two currents formally is the natural next step.

6.2 Statement of Theorem 1L (local)

Theorem 1L (Local Lorentz Emergence). *Under H1–H5 and Proposition T1, at each point $x \in \mathcal{M}$_coarse the structure group of T_x $\mathcal{M}$_coarse compatible with the cone $\mathcal{C}(x)$, transport-density preservation, and time-orientation has connected component $SO^+(1,3)$. The full admissible structure group is at least the orthochronous Lorentz group $O^\uparrow(1,3) = SO^+(1,3) \rtimes \{1, P\}$.*

6.3 Proof of 1L

The proof has four steps.

Step A — Cone to conformal class (Malament; Hawking–King–McCarthy). By the theorem of Malament (1977) and the prior result of Hawking, King, and McCarthy (1976), on a strongly causal, time-orientable Lorentzian manifold of dimension > 2, the causal structure determines the metric up to a conformal factor $\lambda(x) > 0$. Applied to $\mathcal{M}$_coarse with the cone field of Lemma 5.1, this gives a *conformal class* of Lorentzian metrics on each tangent space.

This is the principal v2 correction. Alexandrov–Zeeman characterise causal automorphisms of *flat* Minkowski space — they presuppose Minkowski rather than producing it. Malament–HKM is the correct cone-to-conformal-class theorem for a general Lorentzian manifold.

Step B — Homogeneity and isotropy. H3 (homogeneity) and H4 (isotropy) restrict the conformal factor $\lambda(x)$ to be invariant under spatial translations and rotations of the coarse-grained manifold. In the local tangent-space analysis, this forces λ to be constant on $T_x \mathscr{M}\_coarse$ for fixed x.

Step C — BCB fixes the conformal factor through transport-volume preservation. Malament–HKM (Step A) gives only a conformal class [g]. By Lemma 6.1.d (transport-volume preservation, local form) and Corollary 6.1.e (conformal factor fixed, with global constancy from connectedness and continuity), Ω is a global constant, absorbed into the definition of units and c. The conformal Lorentzian structure is therefore upgraded to a metric Lorentzian structure, and the local structure group reduces from the conformal group to the metric-preserving Lorentz group O(1,3).

Step D — Time-orientation (A2) restricts to O↑(1,3); orientation does not restrict further. A2 (irreversibility) gives a global time-orientation on the cone field, restricting O(1,3) to the orthochronous subgroup O↑(1,3) by excluding time-reversal T. A2 does *not* exclude spatial parity P: parity reflects space and does not reverse the commitment direction. The connected component is therefore $SO^+(1,3)$; the full admissible structure group is at least O↑(1,3) = $SO^+(1,3) \rtimes \{1, P\}$.

This corrects the v1 over-reach. Excluding P would require an independent substrate-level handedness primitive, which the present axiom set does not provide. Physical parity violation in the Standard Model is a downstream *dynamical* breaking story — electroweak chirality — derived elsewhere in the programme rather than imposed kinematically here.

Combining Steps A–D, the connected component of the local structure group is $SO^+(1,3)$; the full admissible structure group is at least O↑(1,3). This completes the proof of Theorem 1L.

6.4 Statement of Theorem 1G (global)

Theorem 1G (Global Poincaré Emergence). *Under the hypotheses of Theorem 1L, if $\mathscr{M}$_coarse admits a flat affine connection compatible with the local $SO^+(1,3)$ structure — i.e., if the substrate produces a flat continuum limit — then the connected component of the global admissible-transformation group is the proper orthochronous Poincaré group $\mathscr{P}^+ = SO^+(1,3) \ltimes \mathbb{R}^{1,3}$.*

6.5 Proof of 1G

If $\mathscr{M}$_coarse is flat, the local structure group of Theorem 1L extends to a global rigid action by parallel transport. On flat Minkowski space, Alexandrov (1950s) and Zeeman (1964) then characterise the group of causal-cone-preserving bijections as the proper orthochronous Poincaré group together with dilations. Proposition T1 removes the dilation subgroup (as in Step C). The connected component is therefore $\mathcal{P}^+$. This completes the proof of Theorem 1G.

6.6 When 1G fails

When the substrate produces a *curved* continuum limit — as in the Friedmann-equation and tensor-perturbation papers — Theorem 1G does *not* apply. What survives is Theorem 1L on each tangent space, together with the diffeomorphism group globally. This is the standard local-Lorentz-structure-plus-Diff($\mathscr{M}$) setting of general relativity. The retrospective justification this paper provides for the curved-substrate papers is therefore *tangent-space-local*, not global Poincaré.

6.7 Why Galilean fails

The Galilean group preserves absolute simultaneity, not transport-cone structure. Under a Galilean boost with parameter v, the future light cone $\mathcal{C}^+(x)$ is mapped to a tilted cone with a different opening angle, violating Lemma 5.1. Galilean is the $c \to \infty$ contraction of Poincaré and is incompatible with finite invariant c (H1).

6.8 Why Euclidean fails

A Euclidean signature metric has no causal cone — every pair of events is "spacelike" in the indefinite sense. This is incompatible with the transport-accessibility partition of Lemma 5.1 and with A2 (irreversibility, requiring a time-orientation).

6.9 Why Carrollian fails

The Carrollian group ($c \to 0$ contraction) collapses the transport cone to a line, making every spatial separation transport-incompatible. This is incompatible with H1 (finite *positive* c) and with the existence of non-trivial commitment propagation.

6.10 Why generic Finsler fails

A Finsler structure with a non-quadratic indicatrix would in general not be cone-preserving under any continuous group larger than the stabiliser of a single point. By the theorem of Deng and Hou (*"The group of isometries of a Finsler space,"* Pacific J. Math. 207, 2002, 149–155) for the positive-definite-signature case, and by the extension of Pfeifer (*"Finsler spacetime geometry in physics,"* Int. J. Geom. Methods Mod. Phys. 16, supp02 (2019), 1941004; arXiv:1903.10185) together with Lämmerzahl and Perlick (*"Finsler geometry as a model for relativistic gravity,"* Int. J. Geom. Methods Mod. Phys. 15, 2018, 1850166) to pseudo-Finsler observer structures in Lorentzian signature, only Riemannian (here pseudo-Riemannian, quadratic) indicatrices admit a 6-dimensional isometry group acting transitively on the *future timelike hyperboloid* (the mass shell). Under H3 and H4, the admissible structure group must be transitive on the future timelike hyperboloid, forcing quadratic structure.

6.11 Why DSR and Hořava–Lifshitz fail structurally

Doubly-special relativity proposes an energy-dependent maximum speed (or equivalently a deformed dispersion relation). Either the deformation introduces an energy-dependent cone — which violates Lemma 5.1 (the cone partition must be observer-invariant and event-pair-defined, with no energy label) — or the deformed symmetries do not form a Lie group, which violates the continuous group assumption underlying the connected-component analysis in Theorems 1L and 1G.

The "relative locality" branch of DSR (Amelino-Camelia–Freidel–Kowalski-Glikman–Smolin) attempts to evade the energy-dependent-cone objection by making locality itself observer-dependent — events that coincide for one observer are spatially separated for another. This is structurally disqualified by Lemma 5.1 in a different way: observer-dependent locality entails observer-dependent cone-partition assignment to event pairs, which is exactly what Lemma 5.1 forbids under H2. The rainbow-gravity variant (energy-dependent metric) is disqualified by Theorem 2 — the invariant interval is uniquely determined up to global units, with no room for energy-dependent rescaling.

Hořava–Lifshitz gravity introduces anisotropic scaling between space and time ($z \neq 1$ Lifshitz exponents). The corresponding "preferred foliation" violates H2 (operational substrate-frame inaccessibility). Anisotropic scaling also violates H4 (isotropy of admissible coarse-graining), since it distinguishes the time direction from spatial directions at the coarse-graining level rather than only at the dynamical level.

Empirical bounds on all branches from high-energy astrophysics (§13) provide an independent constraint.

---

7. Theorem 2 — Derivation of the invariant interval

7.1 Statement

Theorem 2 (Invariant Interval). *Under the hypotheses of Theorem 1L, the unique (up to overall units) quadratic form on $T_x \mathcal{M}_{coarse}$ invariant under $SO^+(1,3)$ is*

$$ds^2 = c^2\, dt^2 - dx^2 - dy^2 - dz^2.$$

7.2 Proof

By Theorem 1L the structure group is $SO^+(1,3)$ on each tangent space, with the conformal factor already pinned to a single global constant by Proposition T1 (Step C of §6.3). The signature follows from the cone being a real, non-degenerate, open convex double cone in 4 dimensions: this forces signature (1, 3) (with the standard convention; the opposite convention is physically equivalent). For uniqueness of the quadratic form up to scale, the relevant fact is that *the space of $SO^+(1,3)$-invariant symmetric bilinear forms on the absolutely irreducible real defining representation is one-dimensional*; equivalently, once $SO^+(1,3)$ is fixed as the local structure group, there exists only one non-degenerate symmetric bilinear form preserved by all admissible transformations, up to overall scale and signature convention. (This is a standard consequence of Schur's lemma; we state the consequence rather than the lemma to avoid the sideways citation in v1.) The one-dimensional invariant is the Minkowski form, with the overall scale fixed by the units of c.

The interval is therefore not primitive geometry. It is the unique bookkeeping invariant of refinement-stable, transport-density-preserving propagation.

---

8. Theorem 3 — Refinement-compatible boost structure

8.1 Refinement compatibility, made precise

Let TPB_k denote the ticks-per-bit layer at refinement level k, and let $\mathcal{R}$: TPB_k → TPB_{k+1} be the refinement functor developed in the sequential-interface-transport and σ-duality papers. A coarse-grained transformation T: $\mathcal{M}$_coarse → $\mathcal{M}$_coarse is *refinement-compatible* iff there exist lifts T_k: TPB_k → TPB_k commuting with $\mathcal{R}$:

$\mathcal{R} \circ T_k = T_{k+1} \circ \mathcal{R}$.

This is H5 made operational.

8.2 The chicken-and-egg resolution

The v1 draft used R1 simultaneously as definition of "boost" and as constraint. The v2 fix: by Theorem 1L, the connected component of the local structure group contains a 3-parameter family of one-parameter subgroups that are neither pure rotations nor (extending to 1G) pure translations. These are the boosts *structurally*. R1 is then the consistency statement that these subgroups lift through $\mathcal{R}$ — a condition that either holds (in which case boosted and unboosted observers see consistent refined-transport limits) or fails (in which case A0 is violated and the framework is inconsistent).

Pending a substrate-level proof in future work, R1 is treated as a labelled conjecture in the present paper. The earlier sequential-interface-transport paper sketches the K=7 wheel-structure intertwiner that is the intended foundation for the proof, pending verification; a full theorem with explicit lift T_k → T_{k+1} commuting with $\mathcal{R}$ is not yet on record. Until that proof exists, Theorem 3 should be read as conditional on R1; the falsification path in §13 (F5) addresses the conjectural case directly.

Theorem 3 should therefore be interpreted not as an independent derivation of Lorentz boosts from first principles, but as the identification of the physically admissible one-parameter subgroup structure once Theorem 1L has fixed the local symmetry group. The non-trivial content is the refinement-compatibility requirement imposed by R1.

8.3 Statement of Theorem 3

Theorem 3 (Boost Generation). *Under H1–H5, Proposition T1, and verification of R1, the one-parameter subgroups of the connected component of G (Theorem 1L / 1G) consisting of refinement-compatible cone- and transport-density-preserving transformations that are not pure rotations or translations are exactly the Lorentz boosts. The boost parameter coincides with the rapidity $\eta \in \mathbb{R}$, related to coarse-grained velocity v by*

$v = c \cdot \tanh \eta$, **with** $\eta = \frac{1}{2} \ln[(c + v)/(c - v)]$.

8.4 Proof sketch

By Theorem 1L (or 1G), the connected component of G is generated by translations, rotations, and boosts. Translations and rotations are accounted for by H3 and H4. The remaining one-parameter subgroups correspond to boosts. R1 ensures these subgroups intertwine with $\boldsymbol{\mathcal{R}}$; this is automatic at the continuum level but is a non-trivial constraint at finite refinement level.

The rapidity-additivity law $\eta_3 = \eta_1 + \eta_2$ follows from the one-parameter-subgroup structure: composition of refinement-compatible transformations is refinement-compatible, and the only one-parameter group structure on the boost subgroup consistent with this and with Lorentz invariance is additive in η.

8.5 Interpretation

Different inertial observers correspond to different refinement-compatible coarse-grained slicings of the same substrate. Boosts are *primarily* transport-equivalence symmetries; they happen to act as Lorentz coordinate transformations on the emergent manifold.

---

9. Emergent Lorentzian geometry

Combining the results so far:

- the cone field is invariant (Lemma 5.1);
- the local structure group is $SO^+(1,3)$ (Theorem 1L); the global symmetry group is $\mathcal{P}^+$ when the continuum limit is flat (Theorem 1G);
- the invariant interval is $ds^2 = c^2\, dt^2 - dx^2 - dy^2 - dz^2$ (Theorem 2);
- boosts are refinement-compatible transport-equivalences with rapidity-additive composition (Theorem 3).

Geometry is downstream of transport, not upstream. The continuum 4-manifold is a coarse-grained encoding of refinement-stable transport accessibility relations.

This justifies retrospectively the *tangent-space-local* Lorentzian structure used in earlier VERSF papers, including the Friedmann-equation and tensor-perturbation papers. It does *not*, on its own, justify their *global curved* manifold structure — that requires the additional dynamical input of those papers (Einstein-equation analogues at the substrate level, expansion dynamics, perturbation evolution) which are not derived here. The honest claim is: each tangent space inherits $SO^+(1,3)$ from this paper; the global curved manifold structure remains the work of the dedicated papers.

This result substantially strengthens the earlier BCB derivations of gauge structure and Standard Model symmetry. In those papers, gauge covariance, spinorial structure, and local comparison of quantum states were developed using Fisher geometry and distinguishability transport, but Lorentzian geometry itself remained inherited. The present paper *removes the obstruction* to closure of that dependency by supplying the local Lorentzian tangent-space structure those derivations require.

What is *not* established here is the derivation-by-derivation re-grounding of the existing BCB gauge papers onto Theorem 1L. That re-grounding requires going back to each BCB gauge paper individually and checking whether its use of Lorentzian structure can be exactly substituted by tangent-space-local $SO^+(1,3)$ (Theorem 1L) or whether it requires global Poincaré covariance (Theorem 1G), which is in turn conditional on the flat continuum limit and does *not* hold in curved-substrate settings. That work is deferred to a dedicated re-grounding paper. The honest claim here is therefore: this paper supplies the local Lorentz structure that the BCB gauge derivations require; whether each specific derivation re-grounds in tangent-space-local form, and which require the additional flat-continuum-limit hypothesis, is the subject of future work.

---

10. Connection to the Dirac-field programme

Earlier VERSF and BCB papers derived:

- spinorial closure structure,
- CAR-algebra emergence,
- Fock-space structure,
- gauge connections from local state comparison,
- U(1) gauge emergence from phase redundancy,
- SU(3) × SU(2) × U(1) structure from projective probability geometry,
- chirality constraints from fold orientation,
- emergent Dirac fields.

Each construction depended at some stage on Lorentz-compatible continuum structure:

- spinors transform under Lorentz representations,
- gauge currents require causal transport structure,
- covariant derivatives require local tangent-space geometry,
- relativistic field propagation requires invariant causal cones.

The present paper supplies that missing geometric layer.

The dependency chain — substrate → cone → $SO^+(1,3)$ → invariant interval → spinor reps → gauge connections → Dirac equation — is now *unblocked* at the tangent-space-local geometric level, in the sense that no step in the chain any longer requires a primitive Lorentzian assumption. Whether the chain *closes* end-to-end depends on (i) the derivation-by-derivation re-grounding of the existing BCB gauge papers onto Theorem 1L (deferred to a dedicated paper, as noted in §9), and (ii) the dynamical content of the Dirac equation, including minimal coupling, which still requires the U(1) and SU(2) × SU(3) gauge structures developed elsewhere in the programme.

---

11. Microcausal structure and the precise open problem

For emergent fermionic field operators $\psi(x)$, $\psi(y)$, full microcausality requires

$\{\psi(x), \psi^\dagger(y)\} = 0$ for spacelike separation $(x - y)^2 < 0$.

This paper *partially* establishes microcausality:

- The transport-cone structure provides the geometric criterion (spacelike separation = transport-incompatibility).
- Transport incompatibility suppresses operational overlap of underlying commitment events.

11.1 What is not yet derived

Full vanishing of the anticommutator requires, in addition, that the emergent field algebra be local in the Haag–Kastler sense — field operators at transport-incompatible events must commute (or anticommute) *exactly*, not merely up to suppressed corrections.

11.2 The precise open problem (v2 sharpening)

The v1 draft flagged a conjectural "absence of substrate-level wormhole admissible-transport paths." We now state this precisely.

Definition (substrate-connectedness). Two coarse-grained events $x, y \in \mathcal{M}$_coarse are *substrate-connected* iff there exists an admissible commitment-transport sequence $(s_0, s_1, \ldots, s_n)$ in $\mathcal{S}$_substrate with $s_0 \in \varphi^{-1}(x)$, $s_n \in \varphi^{-1}(y)$, where φ: $\mathcal{S}$_substrate → $\mathcal{M}$_coarse is the coarse-graining map and each consecutive pair (s_i, s_{i+1}) is locally coupled in the sense of A4.

Conjecture (No-Wormhole Theorem). *Under H1–H5 and BCB, if* $y \notin \mathcal{C}(x)$ *— i.e., y is neither in the future nor in the past transport-cone of x, equivalently x and y are spacelike-separated in the emergent geometry — then x and y are not substrate-connected.*

The cone-based formulation avoids the question of how the interval $ds^2(x, y)$ is defined for non-infinitesimal event pairs — it is stated directly in terms of the cone partition, which is well-defined globally by Lemma 5.1. The conjecture's content is that the coarse-graining map φ does not glue distant substrate regions to coarse-grained event pairs that are cone-disjoint. Algebraic microcausality follows from the conjecture by standard arguments (Reeh–Schlieder-style locality for the emergent field algebra).

11.3 Status

Cone-level microcausality: *proven, conditional on H1–H5 and T1*. Algebraic microcausality: *conditional on the No-Wormhole Conjecture*. The conjecture is now precisely stated and is a candidate for substrate-level proof in future work.

---

12. Consequences for quantum-field emergence

The overall VERSF architecture, with the present paper inserted, is more honestly rendered as a layered DAG with parallel branches rather than a linear chain. The diagram below uses indentation and explicit "←" arrows to indicate which prior results each node depends on; nodes at the same depth are mutually independent or parallel.

```
Layer 1 — Substrate axioms (all parallel):
   A0, A1, A2, A3, A4, BCB

Layer 2 — Immediate consequences (two parallel descent paths):
   information structure                  ← from A0, A1, A2
   finite invariant propagation (H1)      ← from A1, A4

Layer 3 — Information geometry
          (three parallel branches from information structure):
   Fisher geometry                        ┐
   Hilbert structure                      ├── all from information
   projective probability geometry        ┘    structure; mutually
                                               consistent but parallel

Layer 4 — Substrate-frame gauge:
   H2 (hypothesis; defeasible by F1)      ← from H1 + §4 operational
                                            characterisation of admissibility

Layer 5 — Geometric primitives:
   invariant causal cones (Lemma 5.1)     ← from H1, H2
   transport-density preservation         ← from BCB, A0, A1
   (Proposition T1)                         (Lemmas 6.1.a/b/c)

Layer 6 — Main results of this paper:
```

```
   Theorem 1L (local Lorentz emergence)   ← H1–H5 + Prop T1 + Lemma 5.1
        │
        ├── Theorem 1G (global Poincaré)  ← Theorem 1L + flat continuum limit
        ├── Theorem 2  (invariant interval) ← Theorem 1L
        └── Theorem 3  (boost structure)  ← Theorem 1L + R1 (conjecture; see
§8.2)

Layer 7 — Field-theoretic structure
         (two parallel inputs to internal symmetry sectors):
   spinorial representations              ← Theorem 1L (this paper)
                                            + projective probability geometry
   gauge connections from local           ← Fisher geometry + Prop T1
   comparison on Fisher manifolds           + Theorem 1L

Layer 8 — Internal symmetry sectors:
   SU(3) × SU(2) × U(1) structure         ← projective probability geometry
                                            + entropy bounds + anomaly
                                            constraints + gauge connections

Layer 9 — Second quantisation:
   CAR algebra + Fock structure           ← spinorial representations
                                            + Theorem 2

Layer 10 — Dirac equation:
   emergent Dirac fields                  ← CAR algebra + gauge connections
                                            + spinorial representations

Layer 11 — Downstream:
   effective relativistic quantum         ← Layers 7–10
   field structure
   Standard Model reconstruction          ← all of the above + dynamical
   programme                                input not derived here
```

Two structural points are now explicit. First, the BCB-gauge derivations (Layer 7, right branch) and the spinorial-representation derivations (Layer 7, left branch) are parallel inputs to internal-symmetry sectors (Layer 8); neither is derived from the other. Second, gauge connections and spinorial representations both require Theorem 1L from this paper as an input — but the BCB gauge-connection derivation additionally uses Fisher geometry from Layer 3 and transport-density preservation from Layer 5, none of which is derived from Theorem 1L. The dependency graph is therefore not a single chain.

The number of inherited continuum assumptions is significantly reduced. The principal remaining inherited inputs are H3 (homogeneity) and H4 (isotropy) of admissible coarse-graining, both flagged for future substrate-level derivation; flat-continuum-limit availability for Theorem 1G; and the No-Wormhole Conjecture for algebraic microcausality.

---

13. Falsification paths and quantitative bounds

The framework is falsified if any of the following are observed.

(F1) Substrate-frame detection. Observation of any preferred-frame effect not attributable to kinematical effects such as the CMB rest frame. The most systematic empirical bounds come from the Standard-Model Extension (SME) framework of Colladay and Kostelecký, which parametrises Lorentz-violation via coefficients including (a) the $c_{\mu\nu}$ and e_μ coefficients for photon-sector violations and (b) the a_μ, b_μ, $c_{\mu\nu}$, $d_{\mu\nu}$ coefficients for fermion-sector violations. Current bounds compiled in the Kostelecký–Russell Data Tables for Lorentz and CPT Violation (arXiv:0801.0287, updated annually) constrain these coefficients at the 10^{-18}–10^{-20} level (or tighter in current editions) via atomic-clock comparisons, Michelson–Morley-type tests, and astrophysical Lorentz-invariance tests. Detection of non-zero SME coefficients above current thresholds falsifies H2.

(F2) Cone violation. Detection of superluminal propagation of any admissible commitment-bearing signal falsifies H1 and undermines Lemma 5.1.

(F3) Lorentz violation at accessible scales. Energy-dependent dispersion of high-energy photons, neutrinos, or cosmic rays beyond the bounds from Fermi-LAT GRB observations ($\Delta c/c < 10^{-15}$ at GRB energies) or IceCube astrophysical-neutrino constraints would tension the framework and directly disqualify DSR/Hořava-type deformations.

(F4) Continuum-limit failure. If substrate simulation or analytic continuum-limit analysis showed coarse-grained dynamics not converging to Lorentzian structure — for example, exhibiting an emergent Finsler indicatrix with non-quadratic curvature, or anisotropic-scaling continuum behaviour — the framework is falsified.

(F5) Incompatible boost sectors (R1 failure). If different inertial observers were shown to have incompatible refinement-stable transport descriptions, R1 fails and Theorem 3 does not apply. The framework would then need additional structure to recover boost-invariant physics.

(F6) Substrate wormholes. If the No-Wormhole Conjecture (§11.2) were violated by a substrate-level construction, algebraic microcausality fails and the emergent field theory is non-local in the Haag–Kastler sense.

(F7) Failure of transport-volume invariance. If a substrate simulation or empirical analogue showed invariant causal cones but non-invariant committed-record density under admissible observer transformations, then the framework would retain conformal Lorentzian structure (Step A of §6.3) but fail to derive a unique metric. Concretely, a substrate model in which BCB held in integrated form (total bit count $Q(\mathcal{U})$ conserved across Cauchy sections) but the differential transport density transformed inhomogeneously under admissible observers — e.g., position-dependent observer rescalings tolerated by some weakened form of A1 — would exhibit this signature. This would falsify Lemma 6.1.d and Corollary 6.1.e — the conformal-factor-fixing step of Step C — while leaving Lemma 5.1 (cone invariance) intact. The resulting structure would be conformal-Lorentzian rather than metric-Lorentzian.

The theory possesses direct empirical vulnerability at the precision frontier.

---

14. Limitations, open problems, and dependencies

Open problems flagged in this paper.

- *Substrate-level derivation of H3 and H4.* Homogeneity and isotropy are assumed as conditional inputs.
- *Full manifold emergence.* This paper assumes a 4-dimensional connected smooth coarse-grained manifold; dimensional emergence is treated elsewhere.
- *Flat continuum limit for Theorem 1G.* Theorem 1L applies universally; Theorem 1G applies only when the substrate produces a flat continuum limit. The curved-substrate cases retain Theorem 1L locally but require their own dynamical analysis globally.
- *No-Wormhole Conjecture (§11.2).* Necessary for algebraic microcausality.
- *Interacting gauge theories.* Addressed in the BCB–VERSF synthesis and Maxwell-admissibility papers but not unified here.
- *Renormalisation structure.* The admissible coarse-graining functor needs to be related to the RG operator on causal diamonds developed in the admissible-coarse-graining paper.
- *Standard Model reconstruction.* Full SM emergence remains an engineering programme. Parity violation in particular is downstream of the present paper, since §6.3 Step D explicitly concedes P-symmetry at the structural level.

Principal remaining structural challenge. Exact continuum-limit control of the substrate dynamics — the rigorous limit theorem connecting finite-refinement TPB structure to the continuum manifold of this paper — is the largest open mathematical problem.

Important dependency clarification. The strongest remaining mathematical dependency is *not* Lorentz symmetry itself, but the regularity theorem connecting finite substrate transport to a smooth, non-degenerate cone field. The present paper assumes this regularity through H3–H5 and Lemma 5.2. A later continuum-limit paper should prove that the admissible TPB refinement sequence converges to a smooth cone distribution satisfying the causal regularity conditions required by Malament–Hawking–King–McCarthy. Once that is established, the chain from Lemma 5.1 through Lemmas 6.1.a–e to Theorems 1L, 1G, 2, and 3 follows essentially mechanically. The structural emergence of Lorentzian geometry is therefore not where the residual mathematical risk lies; the residual risk lies in the continuum-limit regularity step.

Status summary.

Result	Status
Lemma 5.1 (cone invariance)	Proven, conditional on H1–H2
Lemma 5.2 (cone regularity)	Proven, conditional on H1–H4 + continuum-regularity (cf. sequential-interface-transport paper)
Lemmas 6.1.a–c (current conservation, charge invariance, scale-fixing)	Proven from BCB + A0 + A1
Proposition T1 (transport-density preservation)	Proven from Lemmas 6.1.a–c

Result	Status
Lemma 6.1.d (transport-volume preservation, local form)	Proven from Lemmas 6.1.a–c + T1 + local no-flux-tube construction
Corollary 6.1.e (conformal factor fixed)	Proven from Lemma 6.1.d + connectedness and continuity of Ω on 𝓜_coarse
Theorem 1L (local Lorentz emergence)	Proven, conditional on H1–H5 + T1 + Cor 6.1.e
Theorem 1G (global Poincaré emergence)	Proven, conditional on H1–H5 + T1 + flat continuum limit
Theorem 2 (invariant interval)	Proven, conditional on H1–H5 + T1 + Cor 6.1.e
Theorem 3 (boost generation)	Conjectural (conditional on R1; structurally proven on H1–H5 + T1)
Parity exclusion	*Not* asserted at structural level (downstream dynamical story)
Geometric microcausality	Proven, conditional on H1–H5 + T1
Algebraic microcausality	Conjectural (No-Wormhole Conjecture)
Global form of Lemma 6.1.d	Conditional on no-flux-tube foliation lemma (admissible-coarse-graining paper); not used in this paper, since the local form combined with connectedness suffices for Cor 6.1.e
Substrate-level H3, H4	Open
Continuum-limit regularity theorem	Open (principal remaining mathematical risk)

---

15. Conclusion

Lorentzian geometry emerges from finite distinguishability, irreversible commitment, local transport, observer-invariant content, and refinement-compatible coarse-graining. The cone field is forced by finite propagation and operational substrate-frame inaccessibility (Lemma 5.1). The metric is forced — up to overall units — by transport-density preservation, which is itself a consequence of BCB (Proposition T1). The local structure group is forced by cone-preservation, homogeneity, isotropy, and time-orientation (Theorem 1L). The global symmetry group is the proper orthochronous Poincaré group when the continuum limit is flat (Theorem 1G). The boost structure is forced by refinement compatibility (Theorem 3).

Relativity is therefore not a primitive geometric postulate. It is the unique stable large-scale geometry compatible with finite, observer-invariant transport of irreversible facts. In VERSF, geometry is emergent, causality is transport structure, and the emergent 4-manifold is the coarse-grained bookkeeping of admissible commitment propagation.

The key strengthening over earlier versions is that Lorentzian geometry is *not* obtained from cone structure alone. Cone structure gives only a conformal metric — the equivalence class $[g_{\mu\nu}]$ under $g_{\mu\nu} \to \Omega^2(x)\, g_{\mu\nu}$. The metric itself is fixed by the additional VERSF claim that committed distinctions are conserved substrate records, not observer-dependent bookkeeping artefacts. BCB, together with A1, therefore does more than conserve information: through Lemma 6.1.d (transport-volume preservation) and Corollary 6.1.e (conformal-factor fixing), the joint content of substrate bit-conservation and finite distinguishability fixes the conformal scale of the emergent geometry. This is the step that turns causal order into metric geometry — and it is the step that turns a conformal-Lorentzian framework (which would survive cone invariance alone) into the full Lorentzian framework required for relativistic physics.

The present work also changes the status of the earlier BCB gauge programme. The downward direction is now established: the BCB gauge derivations of gauge connections, spinorial transport, chirality structure, and Standard Model internal symmetry have, in this paper, a substrate-level Lorentzian foundation rather than an inherited one. The upward direction — verifying that the existing BCB gauge derivations are individually consistent with, and re-expressible on, the local-only Theorem 1L (rather than tacitly requiring the stronger global Theorem 1G) — is a separate consistency check that has not been performed in the present paper and is the subject of a dedicated re-grounding paper. The present paper therefore *removes the obstruction* to closure of the programme architecture, in which information-theoretic constraints, gauge structure, Hilbert geometry, causal transport, and emergent relativistic geometry are now part of a single constrained architecture rather than parallel derivation tracks. It does not, on its own, complete that closure.

In particular, to be explicit about scope: this paper does not derive the Standard Model dynamically. What it establishes is that the Lorentzian geometric substrate required by the BCB gauge-construction programme no longer needs to be postulated independently.

The remaining work is to derive H3 and H4 from substrate principles, prove the continuum-limit regularity theorem identified in §14 as the principal remaining mathematical risk, prove the No-Wormhole Conjecture, prove R1 fully at the substrate level, and carry out the derivation-by-derivation re-grounding of the BCB gauge papers onto Theorem 1L identified in §9, §10, and §15.