

Distinguishability, Order, and Universality: A First-Principles Reduction of Relativistic Kinematics in VERSF

Keith Taylor — VERSF Theoretical Physics Programme

General Reader Abstract

Three of the most basic features of modern physics are usually treated as separate starting points. *Time has one dimension*: we move forward through it, not sideways. *The laws of physics look the same to all observers in steady motion*: this is Einstein's principle of relativity. *All objects fall at the same rate in gravity, regardless of what they are made of*: a feather and a hammer dropped on the Moon hit the surface together. (This is what physicists call the *weak* equivalence principle. A stronger version, covering all gravitational effects rather than just free fall, is a separate claim and is not derived here.)

Physics textbooks introduce each of these as its own postulate — a starting assumption you have to accept before the rest of the theory can be built. This paper shows that all three follow from a single, much simpler idea:

Physics can only be about things observers can actually tell apart. Anything an observer cannot detect, even in principle, cannot be part of the physical content of a theory.

Working inside a research programme called VERSF — which builds physics from the ground up out of irreversible "commitment events" rather than from a pre-given spacetime — we apply this single principle to three separate questions. The results are:

- **Time has one dimension** because that is what is left after you strip away every direction along which two observers could disagree about the order of events. Disagreements that *can* exist (between events too far apart for any signal to connect them) turn out to be exactly the spatial directions we already know about.
- **Einstein's relativity emerges in the no-detection branch.** The axioms force a sharp choice: if the substrate's "preferred frame" leaves no measurable fingerprint on real experiments (no committed-record trace), then it is gauge and the Lorentz group follows. If it does leave such a fingerprint, relativity would already have been seen to fail. The world is currently observed to be in the first case to very high precision in many sectors, with experimental bounds tightening progressively.
- **Universal free fall holds** because gravity, in this framework, responds to the substrate of matter — the underlying "stuff" — and not to the labels chemistry and particle physics put on different forms of that stuff.

Each result comes with a clear way to test it experimentally and a clear statement of what would have to be true if the test failed. The paper closes with an explicit accounting of what assumptions remain after the reduction, and exactly how much each has been narrowed.

This paper is a companion to the broader paper *VERSF as the Origin of Special and General Relativity*, which presents the full programme: how, starting from a substrate of irreversible records and finite distinguishability, the structures of special and general relativity emerge as consequences rather than assumptions. That paper covers the whole arc — the substrate axioms, the derivation of the speed limit, three spatial dimensions, the route to gravitational coupling, and the recovery of Einstein's field equations — and identifies three particular reductions that need to be developed in full mathematical rigour: that the Lorentz group is forced rather than assumed, that gravity treats all matter the same way, and that time is one-dimensional.

The present paper is that development. It takes those three reductions and treats each one at full proof-level detail, with sharply named lemmas, an itemised account of every residual assumption, and a falsification path for each result. Where the parent paper shows *that* the three reductions hold, this paper shows exactly how each one is proved, what it depends on, and where the boundary lies between what is derived here and what is imported from elsewhere in the programme. It does not replace the parent paper; it supplies the proof-level work that the parent paper flagged as needing further development.

The technical body of the paper that follows develops these arguments rigorously.

Table of Contents

Abstract (technical)

1. Introduction

2. Foundations 2.1 Axioms (with OC-min minimal formalism) 2.2 Definitions 2.3 The Fold Hypothesis (Part III only)

3. Part I — The Temporal Dimension 3.1 Commitment induces a partial order 3.2 Comparable-pair facthood 3.3 Time function via linear extension 3.4 Main theorem (one-dimensional temporal sector)

4. Part II — Lorentz Invariance 4.1 Finite propagation speed 4.2 Lemma OC-1; Lemma A1 (regime characterisation of F_0) 4.3 Lemma A2: Group structure 4.4 Effective homogeneity, isotropy 4.5 Corollary A1.1: Substrate-frame velocity is gauge 4.6 Corollary A1.2: Causal-cone invariance 4.7 Theorem 3 (Lorentz invariance)

5. Part III — Universal Free Fall (conditional on FH_min) 5.1 Substrate-level versus effective-level distinctions 5.2 Fold Indistinguishability (FH_min) 5.3 Gravity is a substrate-level coupling 5.4 Theorem 4 (Universal coupling)

6. Synthesis: The A0 Gauge Sieve

7. Falsification

8. Dependency Closure and Residual Assumptions D1 Fold Hypothesis (FH_min) D2 Observer-Protocol Formalism (OC-min) D3 Lorentz Regime Selection D4 Spatial Dimension Import D5 Smooth-Manifold Coarse-Graining (CGH.0 decomposed) Status table What this paper does not derive

9. Conclusion

References

Abstract

We give a unified reduction of three structural ingredients of relativistic physics within the Void Energy–Regulated Space Framework (VERSF): the existence and uniqueness of a single temporal dimension, the Lorentz invariance of effective kinematics, and the universality of gravitational free fall (WEP). The reductions share a single epistemic constraint — Observer-Invariant Distinguishability (A0) — combined with substrate axioms (A1 finite distinguishability, A2 irreversible commitment, A4 local coupling), a minimal observer-comparability formalism (OC-min), and a restricted Fold Hypothesis (FH_min) that is the only piece of the upstream fold-density programme this paper actually requires.

The unifying schema is what we call the *A0 Gauge Sieve*: any candidate physical distinction not resolved by admissible observer-protocols is, by A0, demoted to gauge. Its three applications here are: (i) *Time*. The irreversible order $<$ on commitment events admits a real-valued time function τ (Szpilrajn linear extension); admissible τ are unique up to monotone reparametrization on $<$ -comparable pairs and gauge-equivalent on $<$ -incomparable pairs. We define a bespoke quantity — the *temporal rank* of $(\mathcal{E}, <)$, which captures only $<$ -comparable order without constraining incomparable pairs — and prove temporal rank = 1. This quantity is explicitly distinguished from the Dushnik–Miller order dimension. (ii) *Lorentz*. The substrate-frame F_0 admits two empirical regimes: motion relative to F_0 either alters committed records (Lorentz violation, currently undetected) or it does not (NSF holds, Lorentz group follows). The Lorentz claim is therefore: $A0 + A1 + A2 + A4 + OC\text{-min} + (\text{no detected Lorentz violation}) \Rightarrow SO^+(1,3)$. (iii) *WEP*. Conditional on FH_min, gravitational coupling treated as a substrate-level functional forces $\kappa_a = \kappa$.

The paper closes with an explicit Dependency Closure section converting each residual assumption into its minimal sharply-stated form. Each result carries a sharp falsification path.

1. Introduction

Relativistic physics rests on three structural commitments usually adopted as postulates: a 3+1-dimensional spacetime; Lorentz invariance of inertial physics; and the universality of gravitational free fall (WEP). Within VERSF, none of these is taken as primitive.

The core epistemic axiom is:

A0 (Observer-Invariant Distinguishability). Physical content depends only on distinctions that are invariant under all admissible observer-protocols. Non-invariant distinctions are gauge.

A0 functions as a *gauge sieve*: any putative physical structure that admits an observer-dependent representation is, by A0, demoted to gauge. This paper identifies three places where the standard postulates of relativistic physics introduce structure of exactly this kind, and shows that A0 forces the answer in each case.

Two honesty-conditions on the reductions are made explicit at the outset and developed fully in §8, alongside one statement of what *is* fully delivered:

- *Part I is unconditional under A0 + A2 + OC-min.* The temporal-sector result (Theorem 1) is derived without empirical regime selection or upstream content imports.
- *Lorentz is conditional on regime selection.* The axiomatic content is a regime characterisation: either Lorentz invariance is empirically broken, or NSF holds and the Lorentz group follows. The Lorentz claim is the second branch, conditional on the empirical fact that no Lorentz violation has been detected.
- *Part III is conditional on FH_min.* The minimal Fold Hypothesis we require — that bound matter carries no primitive substrate labels beyond commitment density — is upstream content from the K=7 / structural-uniqueness companion paper. Part III's WEP derivation is therefore a conditional reduction, but it is conditional on the *minimal* form of FH, not the full fold-density programme.

All three statuses are made structurally explicit in the relevant theorems and accounted for in §8.

2. Foundations

2.1 Axioms

We adopt the following axioms of the VERSF substrate, taken as upstream:

- **A0 (Observer-Invariant Distinguishability).** Physical content depends only on distinctions invariant under all admissible observer-protocols. Non-invariant distinctions are gauge.
- **A1 (Finite Distinguishability).** Any bounded substrate region admits only finitely many operationally distinguishable states per unit interval.
- **A2 (Irreversible Commitment).** The substrate supports irreversible "commitment" events; reversal incurs a finite thermodynamic cost. Sub-commitment substrate updates are not records.
- **A4 (Local Coupling).** Substrate dynamics couple states only within a bounded neighbourhood per update.

The observer-protocol structure is given by a minimal formalism, OC-min, sufficient for all results of this paper:

OC-min (Minimal Observer-Protocol Formalism). An admissible observer-protocol is a map $\mathcal{O} : \mathcal{R}^k \rightarrow \mathbb{R}$, where \mathcal{R} is the set of A2-committed records, satisfying:

1. **Finiteness.** \mathcal{O} uses finitely many records.
2. **Commitment access.** \mathcal{O} acts only on A2-committed records and comparisons among them; it does not access pre-commitment substrate updates.
3. **Local constructibility.** \mathcal{O} is built from A4-local comparison primitives.
4. **Gauge respect.** If two inputs are A0-equivalent, \mathcal{O} produces the same output.
5. **Closure under composition.** Finite compositions of admissible protocols are admissible.

We denote the class of such protocols \mathcal{O} .

A4 is used essentially in Theorem 2 (finite signal speed); OC-min Clause 2 is the load-bearing premise of Lemmas OC-1 and A1 below. The clean dependency lists are recorded with each theorem statement.

2.2 Definitions

Let $\mathcal{A}_{\mathcal{O}}$ be the algebra generated by the outputs of admissible protocols.

Definition 1 (Observable equivalence). States x, y are *observationally equivalent* if $\mathcal{O}(x) = \mathcal{O}(y)$ for every $\mathcal{O} \in \mathcal{O}$. Physical laws are defined on equivalence classes.

Definition 2 (Commitment event). A commitment event e is an A2-irreversible substrate update. Let \mathcal{E} denote the set of commitment events.

Definition 3 (Fact). A *fact* is a commitment event $e \in \mathcal{E}$ whose content, together with all order-relations *and order-non-relations* between e and other commitments, is invariant across all $\mathcal{O} \in \mathcal{O}$.

The non-relation clause is essential: two $<$ -incomparable commitments are still required to be incomparable in every admissible representation, so that "spacelike separation" — which exists in relativistic theories — is an invariant fact, not an artefact of representation choice.

Definition 4 (Operational detectability). A structure \mathcal{S} attached to states is *operationally detectable* iff there exist $\mathcal{O} \in \mathcal{O}$ and configurations x, x' such that $\mathcal{O}(x; \mathcal{S}) \neq \mathcal{O}(x'; \mathcal{S})$.

Definition 5 (A0 gauge sieve). A candidate physical structure \mathcal{S} *survives the A0 gauge sieve* iff every distinction it makes is invariant under \mathcal{O} . The non-invariant component of \mathcal{S} is gauge.

2.3 The Fold Hypothesis (Part III only)

For Part III we require only the minimal restriction of the Fold Hypothesis:

(FH_min) Restricted Fold Hypothesis. Bound matter carries no primitive substrate labels beyond commitment density. Equivalently, no substrate-level observer-protocol can distinguish matter types except through total bound-record density.

The full Fold Hypothesis of the upstream programme — that all bound matter is composed of configurations of a single primitive fold — is sufficient for FH_min but stronger than we need here. The present paper proves only that $A0 + FH_min \Rightarrow WEP$; the derivation of FH_min from the K=7 closure architecture belongs to the upstream fold-density and structural-uniqueness papers. This narrowing is recorded as D1 in §8.

3. Part I — The Temporal Dimension

3.1 Commitment induces a partial order

Proposition 1 (Order from commitment). A2 induces a relation $< \subset \mathcal{E} \times \mathcal{E}$ such that $(\mathcal{E}, <)$ is a strict partial order: irreflexive, transitive, and acyclic.

Proof. Define $e_i < e_j$ iff e_i causally enables e_j . Irreflexivity: an event cannot enable itself within a single commitment. Transitivity: chained enablement. Acyclicity: a cycle would imply zero-cost reversal, contradicting A2. ■

The order is in general not total: $<$ -incomparable pairs are permitted and represent commitments with no causal-enablement relation in either direction.

3.2 Comparable-pair facthood

Proposition 2 (Comparable-pair invariance). If $e_i < e_j$ is a fact in the sense of Definition 3, then this relation is invariant across all $\mathcal{O} \in \mathcal{O}$.

Proof. If two protocols $\mathcal{O}_1, \mathcal{O}_2$ disagreed, " $e_i < e_j$ " would depend on the protocol; by Definition 3 it would not be a fact, contradicting the hypothesis. ■

The dual statement applies to non-relations: if e_i, e_j are $<$ -incomparable as a fact, all admissible protocols must record this incomparability. The order of incomparable pairs is not constrained — different admissible protocols may impose different (gauge) orderings on incomparable pairs without violating Definition 3.

3.3 Time function via linear extension

Theorem (Szpilrajn, 1930). Every partial order on a set extends to a linear (total) order on the same set.

For $(\mathcal{E}, <)$ we therefore have a (non-empty) family $\mathcal{L} = \{\preceq\}$ of linear extensions. Each $\preceq \in \mathcal{L}$ refines $<$: if $e_i < e_j$ then $e_i \preceq e_j$; for $<$ -incomparable pairs, \preceq supplies an order absent in $<$.

Definition 7 (Time function). A *time function* is a map $\tau : \mathcal{E} \rightarrow \mathbb{R}$ such that $e_i < e_j \implies \tau(e_i) < \tau(e_j)$. Equivalently, τ is an admissible time function iff it is a real-valued labelling of some linear extension $\preceq \in \mathcal{L}$ — i.e., there exists $\preceq \in \mathcal{L}$ such that for all $e_i, e_j \in \mathcal{E}$, $\tau(e_i) < \tau(e_j)$ iff $e_i \preceq e_j$.

The equivalence is essential for the gauge claim of part (c) below: it pins each admissible τ to a specific linear extension, so that the (c) statement "different linear extensions give different time functions, and the difference is gauge" has a clean target.

3.4 Main theorem (one-dimensional temporal sector)

The standard order-theoretic measure of poset complexity is the *Dushnik–Miller order dimension*: the minimum k such that the poset embeds into \mathbb{R}^k with the product order, faithfully encoding *both* relations and non-relations. For typical non-total posets this can be ≥ 2 and is generically not 1.

This is *not* the quantity we want. The physical content we are after is the rank needed to encode $<$ -comparable order — non-relations on incomparable pairs are gauge by Proposition 2's dual and need not be encoded. We therefore introduce a bespoke quantity tailored to this purpose.

Definition 7' (Temporal rank). The *temporal rank* of $(\mathcal{E}, <)$ is the minimum k such that there exist real-valued functions $\tau_1, \dots, \tau_k : \mathcal{E} \rightarrow \mathbb{R}$ whose joint values determine $<$ -comparable order on every comparable pair. No constraint is placed on the joint values for $<$ -incomparable pairs.

Note. Temporal rank is not the Dushnik–Miller order dimension. Order dimension requires the joint values to determine the *full* product order, including the absence of relations on incomparable pairs; temporal rank requires only that $<$ -comparable pairs are correctly ordered. For typical posets, order dimension \geq temporal rank, often strictly. The argument below establishes temporal rank = 1, which is the quantity physically relevant to the temporal sector; it does not, and does not need to, establish that the order dimension of $(\mathcal{E}, <)$ is 1.

Theorem 1 (One-dimensional temporal sector from facthood). Under A0, A2, and OC-min:

- (a) A time function $\tau : \mathcal{E} \rightarrow \mathbb{R}$ exists (by Szpilrajn).
- (b) On \prec -comparable pairs, the order recorded by τ is uniquely determined by \prec and is invariant across admissible protocols (Proposition 2).
- (c) On \prec -incomparable pairs, the order recorded by τ depends on the choice of linear extension $\preceq \in \mathcal{L}$; by A0, this dependence is gauge.
- (d) The temporal rank of (\mathcal{E}, \prec) is exactly 1. The temporal sector is therefore a one-parameter (real-valued) labelling, fixed up to monotone reparametrization on chains and gauge transformations on \prec -incomparable pairs.

The temporal dimension is exactly one.

Proof.

- (a) Szpilrajn extension theorem applied to (\mathcal{E}, \prec) yields $\preceq \in \mathcal{L}$; any monotone real-valued labelling of (\mathcal{E}, \preceq) gives τ .
- (b) If two admissible τ, τ' disagreed on the order of a comparable pair, then " $e_i < e_j$ " would be protocol-dependent, contradicting Proposition 2.
- (c) Different linear extensions $\preceq, \preceq' \in \mathcal{L}$ generically order \prec -incomparable pairs differently. By Definition 7, the corresponding τ, τ' disagree on those pairs. By Definition 5 and A0, this disagreement is gauge. The dual non-relation clause of Definition 3 is preserved: both τ and τ' record that the underlying \prec -relation is absent.
- (d) *Upper bound (temporal rank ≤ 1).* By (a), there exists a time function $\tau : \mathcal{E} \rightarrow \mathbb{R}$ such that for any \prec -comparable pair $e_i < e_j$: $\tau(e_i) < \tau(e_j)$. Hence the single function τ alone determines \prec -comparable order on every comparable pair. By Definition 7', temporal rank ≤ 1 .

Lower bound (temporal rank ≥ 1). By A2 there exist commitment events related by causal enablement, hence at least one \prec -comparable pair. The empty tuple of functions records no order; a single non-empty tuple is the minimum needed to record a nontrivial pairwise order. Hence temporal rank ≥ 1 .

Therefore temporal rank = 1. ■

Corollary 1.1 (Arrow of time). Along any chain, τ is monotonically increasing, with orientation fixed up to a global sign by A2.

Corollary 1.2 (Spatial sector as gauge quotient on \prec -incomparable pairs). The structure carried by \prec -incomparable pairs is independent of the temporal sector and is modded out by the A0 gauge action on linear extensions. The dimensionality of this quotient — the spatial sector —

is determined by separate stability and conservation arguments (the upstream "n_space = 3" result imported in §8 D4), not by Theorem 1.

Remark 1.3 (Relativistic compatibility). Spacelike-separated events in standard relativity are \prec -incomparable. Their \prec -non-relation is invariant by Definition 3; their order in any frame's time coordinate is gauge by part (c). This is exactly what relativistic theories assert about spacelike separation.

Remark 1.4 (Why temporal rank, not order dimension). A reader familiar with order theory might expect a "dimension" claim about (\mathcal{E}, \prec) to be a Dushnik–Miller order-dimension claim, and might object that the order dimension of a generic causal poset is ≥ 2 . That objection would be correct about order dimension but does not bear on Theorem 1. The temporal sector is defined to be the structure that encodes \prec -comparable order, with incomparability gauge-quotiented. Temporal rank, defined for exactly that purpose, equals 1.

4. Part II — Lorentz Invariance

This part is structured as an explicit chain: Theorem 2 (finite speed) \rightarrow Lemma OC-1 (no pre-commitment access) \rightarrow Lemma A1 (regime characterisation of F_0) \rightarrow Lemma A2 (group structure) \rightarrow effective homogeneity and isotropy \rightarrow Corollaries A1.1, A1.2 \rightarrow Theorem 3 (Lorentz invariance via Ignatowski).

4.1 Finite propagation speed

Theorem 2 (Finite signal speed). Under A1 and A4 there exists a finite $c < \infty$ such that no admissible observer-protocol detects record propagation faster than c .

Proof. A4 bounds the spatial reach of substrate updates per A1-tick by a finite radius ℓ . A1 bounds the number of distinguishable updates per unit duration by a finite rate v . Hence record-relevant propagation is bounded by $c := \ell v < \infty$. ■

(c identified here is the same c that appears in the upstream closure-scale paper, where it equals $\xi \times$ tick-rate.)

4.2 Lemma OC-1 and Lemma A1

Lemma OC-1 (No pre-commitment access). Under OC-min, any substrate distinction that leaves no trace in committed records is undetectable by any admissible observer-protocol.

Proof. By OC-min Clause 2, admissible protocols access only A2-committed records and their comparisons. A substrate distinction with no committed-record trace yields identical outputs for all such protocols. Hence the distinction is operationally inaccessible. ■

This is exactly the structural fact that closes the standing criticism "*maybe an observer can access the substrate frame directly*": OC-min Clause 2 forbids it by definition.

Let F_0 denote a putative *substrate update frame*: the frame in which elementary substrate updates are synchronised before coarse-graining.

Lemma A1 (Substrate-frame regime characterisation). Under A0, A1, A2, and OC-min, the operational status of F_0 is governed by which of the following two regimes obtains in the empirical world:

(a) Substrate-frame motion alters committed records, in which case F_0 is empirically detectable, propagation is observably anisotropic, and Lorentz invariance is empirically falsified.

(b) Substrate-frame motion does not alter committed records, in which case (by Lemma OC-1) no admissible observer-protocol detects F_0 , the substrate-frame distinction is gauge by A0, and NSF holds.

Current experimental bounds strongly support treating the observed world as regime (b) at tested sensitivities.

Proof. Suppose $\mathcal{O} \in \mathcal{O}$ distinguishes two systems S_1, S_2 identical in all committed-record content except for uniform velocity relative to F_0 . By OC-min Clause 2, the distinction is encoded in some difference $\Delta\mathcal{R} \neq 0$ in committed records.

The substrate frame is defined by synchronisation of pre-commitment substrate updates. By A2, these updates are not records. By Lemma OC-1, sub-commitment substrate distinctions that do not alter committed records cannot be detected by any admissible \mathcal{O} .

Two regimes are therefore possible:

(a) F_0 -motion *does* alter committed records. Then $\Delta\mathcal{R} \neq 0$ has an empirical signature — anisotropic propagation, direction-dependent timings, or a measurable preferred frame. In this regime Lorentz invariance is empirically violated.

(b) F_0 -motion does *not* alter committed records. Then by Lemma OC-1, no admissible \mathcal{O} distinguishes S_1 from S_2 ; the supposition fails; by A0, F_0 -velocity is gauge.

Current experimental bounds support treating the world as regime (b). ■

Remark A1.1 (Honest framing). The lemma is not a derivation of regime (b) from the axioms; it is a *characterisation* of which axiomatic consequences obtain in each empirical regime. The selection of branch (b) is an empirical fact. This is a feature, not a weakness: the falsification path of the theory is precisely regime (a). The mature framing is recorded under D3 in §8.

4.3 Lemma A2: Group structure of admissible inertial transformations

Lemma A2 (Group closure). The set G of transformations between admissible inertial observer-protocols is a group: it contains the identity, is closed under composition, contains inverses, and is associative.

Proof. (i) Identity: the trivial transformation between an observer-protocol and itself is admissible. (ii) Composition: if $\mathcal{O}_1 \rightarrow \mathcal{O}_2$ is admissible and $\mathcal{O}_2 \rightarrow \mathcal{O}_3$ is admissible, then the composed protocol-relation $\mathcal{O}_1 \rightarrow \mathcal{O}_3$ is the composition of two A2-grounded, A4-local procedures, which is itself admissible by OC-min Clause 5. (iii) Inverses: each admissible transformation is invertible because A0-equivalence is symmetric. (iv) Associativity: composition of finite procedures is associative. ■

4.4 Effective homogeneity, isotropy

By Lemma OC-1 reasoning applied to spatial location and direction: any preferred location or direction at the substrate level contributes to physical content only if operationally detectable through committed records. Where it is not so detected, A0 makes it gauge. Hence within the admissible regime (b), the kinematics is effectively homogeneous and isotropic.

Effective-level preferred frames or directions — for example, the cosmic microwave background rest frame, or any thermal-state-defined frame — are not at issue here. Those are state-level features (a particular distribution of records in this universe), not substrate-level features of the kinematic group. The hom/iso claim above concerns what features can affect the inertial transformation group between admissible observer-protocols, not what features can exist in any particular configuration. A universe with a CMB rest frame and exact Lorentz invariance is consistent: the CMB picks out a frame for thermal physics without making the kinematic group depend on it.

This is the same regime conditional as in Lemma A1: substrate-level location/direction features that *do* alter committed records constitute their own empirical signatures and would be detected. Their absence at any tested level places us in the homogeneous-isotropic regime.

4.5 Corollary A1.1: Substrate-frame velocity is gauge

Corollary A1.1. In regime (b), $v_0(S)$ — the velocity of system S relative to F_0 — is a gauge parameter, not a physical observable.

Proof. By A0, physical content is the equivalence class $[S] = \{ S' : \mathcal{O}(S') = \mathcal{O}(S) \text{ for all } \mathcal{O} \in \mathcal{O} \}$. By Lemma A1(b), configurations differing only in v_0 lie in the same equivalence class. Hence v_0 does not label physical content. ■

4.6 Corollary A1.2: Causal-cone invariance

Corollary A1.2. Given Theorem 2, Corollary A1.1, and effective homogeneity/isotropy (§4.4), all admissible inertial observer-protocols agree on the same causal boundary:

$$\|\mathbf{x}\| = c t.$$

Proof. Suppose two admissible inertial observers F, F' assigned different values c, c' to the maximum record-propagation speed.

Step 1. By effective homogeneity and isotropy (§4.4), the value $c(F)$ assigned by frame F cannot depend on position (homogeneity) or direction (isotropy).

Step 2. Hence any frame-dependence of c must arise from differences between inertial frames themselves.

Step 3. Under homogeneity and isotropy, the only admissible distinction between inertial frames is their relative velocity. No position label or direction label is available.

Step 4. If $c(F)$ differed between frames as an intrinsic property of a single frame rather than a symmetric pairwise relation, that distinction would require a background reference frame. The only candidate background reference is F_0 . Hence frame-dependent c would amount to tracking velocity relative to F_0 , contradicting Corollary A1.1.

Therefore $c' = c$, and the causal boundary is invariant. ■

4.7 Theorem 3 (Lorentz invariance)

Theorem 3 (Lorentz invariance). Given:

- finite invariant causal speed c (Theorem 2);
- substrate-frame inaccessibility in regime (b) (Lemma A1, Corollaries A1.1–A1.2);
- effective homogeneity and isotropy (§4.4);
- group closure of admissible inertial transformations (Lemma A2);
- *empirical hypothesis: current experimental bounds support treating the world as regime (b) at tested sensitivities, with no Lorentz violation detected;*

the admissible transformation group is the Lorentz group $SO^+(1,3)$, and the Minkowski quadratic form

$$ds^2 = c^2 dt^2 - dx^2 - dy^2 - dz^2$$

is invariant.

Proof. The premises are precisely the inputs of the von Ignatowski derivation. Effective homogeneity gives linearity. Isotropy and group closure (Lemma A2) yield the standard one-parameter classification by an invariant K , with three cases:

$K = 0$: Galilean group; excluded by causal-cone invariance. (The Galilean limit $c \rightarrow \infty$ degenerates the timelike cone to a foliation by simultaneity surfaces, which is not invariant under finite- c boosts; equivalently, Galilean transformations do not preserve $\|\mathbf{x}\| = c t$ for finite c .)

$K < 0$: Imaginary speed; excluded by Theorem 2.

$K > 0$: Lorentz group with $c^2 = 1/K$. Admissible.

Hence $G = \text{SO}^+(1,3)$, with the Minkowski form invariant. ■

Remark 3.1 (Status of CGH.1). The frame-invariance assumption CGH.1 in the upstream VERSF axiom set is now reduced to a derived consequence of $A0 + A1 + A2 + A4 + \text{OC-min}$ plus the empirical regime-(b) selection.

Remark 3.2 (What is not derived). Theorem 3 gives the kinematic group and invariant quadratic form. It does not derive global manifold structure or curvature dynamics; those are downstream of additional VERSF inputs. The remaining manifold burden is decomposed and tracked under D5 in §8.

5. Part III — Universal Free Fall (conditional on FH_min)

Part III is a conditional reduction: WEP follows from $A0 + \text{FH_min}$, where FH_min is the minimal restriction of the Fold Hypothesis stated in §2.3.

5.1 Substrate-level versus effective-level distinctions

A subtlety worth making explicit: composition *is* operationally distinguishable. Chemistry distinguishes gold from copper. The puzzle of WEP is not "matter has no labels"; it is "why does gravity refuse to use the labels matter has?"

We resolve this by a level distinction:

- **Substrate-level quantities** are those resolvable by primitive substrate operations: commitment density ρ_{bound} , fold configurations, and the geometric structures derived from them.
- **Effective-level quantities** are derived classifications produced by composite observer-protocols (chemistry, spectroscopy, scattering experiments).

FH_min asserts that the substrate level admits no primitive type label beyond commitment density.

5.2 Fold Indistinguishability (FH_min implementation)

Lemma 5 (Fold Indistinguishability). Under FH_min , any two configurations $\mathcal{C}_1, \mathcal{C}_2$ with equal bound-record density are *substrate-level indistinguishable*: no admissible observer-protocol acting at the substrate level resolves their difference,

$$\rho_{\text{bound}}(\mathcal{C}_1; x) = \rho_{\text{bound}}(\mathcal{C}_2; x) \implies \mathcal{O}_{\text{sub}}(\mathcal{C}_1) = \mathcal{O}_{\text{sub}}(\mathcal{C}_2) \text{ for all } \mathcal{O}_{\text{sub}} \in \mathcal{O}_{\text{sub}},$$

where $\mathcal{O}_{\text{sub}} \subset \mathcal{O}$ is the subclass of admissible protocols whose outputs depend only on substrate-level quantities.

Proof. By FH_min, the only substrate-level observable is commitment density. A protocol in \mathcal{O}_{sub} depends only on substrate-level quantities; two configurations with equal ρ_{bound} therefore yield equal output. ■

Lemma 5 is the implementation of FH_min that the WEP proof actually uses. The chain is:

FH_min \Rightarrow Lemma 5 (Fold Indistinguishability) \Rightarrow Theorem 4 ($\kappa_a = \kappa$) \Rightarrow WEP.

Lemma 5 does *not* claim that all admissible protocols (effective-level included) fail to distinguish \mathcal{C}_1 from \mathcal{C}_2 ; chemistry succeeds. The claim is restricted to substrate-level protocols.

5.3 Gravity is a substrate-level coupling

In the VERSF programme, gravitational loading is by construction a substrate-capacity response. Capacity is counted at the fold level, not at the level of emergent type-labels. Hence the gravitational loading functional Σ belongs to \mathcal{O}_{sub} .

This is the structural premise of WEP in VERSF. It is not an additional assumption beyond FH_min; it is a statement about which physical quantity gravity is.

5.4 Theorem 4 (Universal coupling, conditional on FH_min)

Theorem 4 (Universal coupling). Given A0, FH_min, and that gravitational loading $\Sigma \in \mathcal{O}_{\text{sub}}$ is a substrate-level functional, then in any effective-level decomposition

$$\Sigma(x) = \sum_a \kappa_a \rho_a(x),$$

the loading coefficients are equal: $\kappa_a = \kappa$ for all a . The gravitational coupling is universal.

Proof. Consider configurations $\mathcal{C}_1, \mathcal{C}_2$ at common spacetime point x with equal total bound density but differing effective composition:

$$\sum_a \rho_a(x) = \sum_a \rho'_a(x), \rho_a(x) \neq \rho'_a(x) \text{ for some } a.$$

Step 1. By Lemma 5, \mathcal{C}_1 and \mathcal{C}_2 are substrate-level indistinguishable.

Step 2. Since $\Sigma \in \mathcal{O}_{\text{sub}}$, $\Sigma(\mathcal{C}_1; x) = \Sigma(\mathcal{C}_2; x)$.

Step 3. Suppose for contradiction $\kappa_a \neq \kappa_b$. Choose $\mathcal{C}_1, \mathcal{C}_2$ with $\rho_a - \rho'_a = \rho'_b - \rho_b \neq 0$ (consistent with equal ρ_{bound}). Then

$$\Sigma(\mathcal{C}_1; x) - \Sigma(\mathcal{C}_2; x) = (\kappa_a - \kappa_b)(\rho_a - \rho'_a) \neq 0,$$

contradicting Step 2.

Step 4 (A0-detectability). A non-universal Σ would induce different gravitational accelerations $\mathbf{a}_1 \neq \mathbf{a}_2$ for substrate-level-indistinguishable configurations. These are operationally detectable (Eötvös-type protocols suffice), so the candidate Σ promotes a non-substrate-level distinction to physical content with substrate-level consequences. By A0 — applied to Σ as a substrate-level coupling — this is forbidden.

Step 5. $\kappa_a = \kappa$ for all a , and $\Sigma(x) = \kappa \rho_{\text{bound}}(x)$. ■

Corollary 4.1 (WEP). The gravitational acceleration of a test body is

$$\mathbf{a} = -\nabla\Phi,$$

independent of effective composition.

6. Synthesis: The A0 Gauge Sieve

The three results share a common logical form:

1. Identify two configurations distinguishable by $\mathcal{S}_{\text{cand}}$ but indistinguishable at the relevant invariance level.
2. Show the resulting distinction is operationally detectable (or operationally inaccessible, in the OC-min sense).
3. By A0, collapse the gauge component.

Domain	Candidate $\mathcal{S}_{\text{cand}}$	Witness	Admissible quotient
Time	Encoding of $(\mathcal{E}, <)$ by k real-valued functions	(i) $<$ -comparable order: invariant by Prop 2; (ii) incomparable-pair order: gauge by A0 on linear extensions	Temporal rank = 1 (= \mathbb{R}^1 labelling)
Inertial frames	Substrate frame F_0 with detectable motion	Lemma A1 regime characterisation; regime (b) selected empirically	$SO^+(1,3)$
Matter coupling	Type-decomposed loading $\{\kappa_a\}$	Configs with equal ρ_{bound} , different ρ_a	Universal κ

(Time, Lorentz, WEP) = $(n_{\text{time}} = 1, \Lambda \in SO^+(1,3)$ [regime (b)], $\kappa_a = \kappa$ [given FH_min]).

Combined with the imported spatial admissibility result (D4 in §8):

$D = 3 + 1, \Lambda \in SO^+(1,3), \text{WEP universal.}$

7. Falsification

Each theorem carries a sharp falsification path.

Theorem 1 (Time). The theorem predicts that for any $<$ -comparable pair, every admissible recording protocol records the pair in the same order — Proposition 2's content. The falsification mode is *structural rather than experimental*: a pair of causally related ($<$ -comparable) commitment events whose order is observer-dependent under all admissible recalibrations and frame choices would mean that $<$ itself is not a fact-producing relation on that pair, contradicting Proposition 2 and the underlying A2 + OC-min structure. (Note that for $<$ -comparable events in standard relativity, all inertial frames already agree on order, so there is no relativistic ambiguity to correct away — the falsification mode is genuinely structural, asking whether $<$ is well-defined as a fact-producing relation, not whether clocks synchronise.) No such structural failure has been identified.

Lemma A1 / Theorem 3 (Lorentz). Lemma A1's regime-(a) branch *is* the falsification path. Any committed-record signature of substrate-frame motion — anisotropic propagation, direction-dependent inertial physics, a measurable preferred frame — would place us in regime (a) and falsify the Lorentz reduction. Current bounds reach $\sim 10^{-20}$ for some SME coefficients, with sensitivities varying widely across the coefficient table; see Kostelecký–Russell for current values. The falsification claim is therefore: "the theory predicts the regime characterisation; experiment selects which branch."

Theorem 4 (WEP). Composition-dependent free fall above the experimental bound (currently $|\eta| < \sim 10^{-15}$, MICROSCOPE final) would falsify FH_min: it would require a substrate-level observable beyond fold-configuration count, contradicting the K=7 no-go result. A confirmed Eötvös-class anomaly would therefore be a *structural* falsification of the upstream uniqueness theorem, not merely a parameter adjustment.

In each case the predicted *value* and the *failure mode* are explicit.

8. Dependency Closure and Residual Assumptions

The preceding results reduce the standard kinematic inputs of relativity, but they do not eliminate every dependency. For clarity, we separate the residual dependencies, state exactly how each is treated in this paper, and record the net change in status produced by the reductions above. Each item is a "closure" in the sense that the dependency is converted from an open-ended assumption into a sharply stated minimal condition.

D1. Fold Hypothesis

The WEP derivation uses only the restricted Fold Hypothesis:

FH_min: bound matter has no primitive substrate labels beyond commitment density.

The full derivation of FH_min from the K=7 closure architecture belongs to the fold-density and structural-uniqueness papers. The present paper proves only

$A0 + FH_min \Rightarrow \text{Lemma 5 (Fold Indistinguishability)} \Rightarrow \text{Theorem 4} \Rightarrow \text{WEP}$.

WEP is therefore not assumed directly; it is reduced to a sharply stated substrate-label condition.

D2. Observer-Protocol Formalism

The paper requires only the minimal observer-protocol structure OC-min defined in §2.1: an admissible observer-protocol is a finite, A4-local procedure acting on A2-committed records and their comparisons, closed under composition.

By Lemma OC-1, any substrate distinction that leaves no committed-record trace is invisible to all admissible protocols. This is the only OC content used in NSF (Lemma A1), in the temporal-sector reduction (Theorem 1), and in the substrate-level/effective-level distinction underlying Theorem 4. Full categorical formalisation of \mathcal{O} — with admissibility measures, monoidal structure, and explicit construction of composite protocols — is deferred to a companion paper, but is not required for any result here.

D3. Lorentz Regime Selection

The Lorentz result is a regime characterisation:

substrate-frame traces in committed records \Rightarrow Lorentz violation detectable,

or:

no such traces \Rightarrow substrate frame is A0-gauge $\Rightarrow SO^+(1,3)$.

Schematically:

$A0 + A1 + A2 + A4 + \text{OC-min} \Rightarrow \{ \text{regime (a): Lorentz violation detectable; regime (b): NSF and Lorentz group} \}$.

VERSF does not assert Lorentz invariance by fiat. It predicts the regime characterisation. Present observations select branch (b) to current experimental precision; this empirical input is what converts the characterisation into the standard relativistic result. The framing turns the residual conditional into the theory's falsification path.

D4. Spatial Dimension Import

The 3+1 decomposition combines this paper's temporal result with an upstream spatial result imported from the spatial-sector reductions of the VERSF programme. The form imported here is:

Spatial Admissibility (imported result). *Under BCB conservation, finite-capacity stability, quantum bound-state stability, and long-range regularity, the only admissible spatial dimension is $n_{space} = 3$.*

The argument structure (developed in the upstream spatial-sector work): $n = 1, 2$ fail long-range and separability conditions; $n \geq 4$ fail classical or quantum bound-state stability; therefore $n = 3$. The detailed proof belongs to the upstream paper(s); the present paper imports the result and supplies the complementary temporal and Lorentz components.

The temporal, Lorentz, and WEP results of this paper are independent of the precise route by which $n_{space} = 3$ is established. They require only that *some* such argument is available; alternative routes to $n_{space} = 3$ (e.g. via stability of three-body Coulomb-like systems, anthropic-from-stability arguments, or direct topological constraints) would suffice equally well. Should the upstream spatial proof require revision, only the 3+1 combination claim of this paper is affected; the temporal-sector reduction (Theorem 1), the regime characterisation (Lemma A1), and the WEP reduction (Theorem 4) stand independently.

Combined:

$n_{time} = 1$ (Theorem 1, this paper), $n_{space} = 3$ (Spatial Admissibility, imported), $D = 3 + 1$.

D5. Smooth-Manifold Coarse-Graining

Even with dimension count and Lorentz group derived, the standard Continuum-Geometry Hypothesis (CGH.0) of the relativity literature contains additional content. We decompose:

$CGH.0 = CGH_{top} + CGH_{smooth} + CGH_{metric}$,

where:

- **CGH_metric (causal/Lorentzian metric structure).** *Derived* in regime (b) of this paper: $SO^+(1,3)$ and $ds^2 = c^2 dt^2 - dx^2 - dy^2 - dz^2$.
- **CGH_top (manifold topology).** Imported from closure / coarse-graining considerations (closure scale ξ , K=7 architecture).
- **CGH_smooth (differentiability).** A continuum-limit hypothesis: finite substrate \rightarrow smooth effective manifold. *Reduction direction (in principle):* the closure scale ξ provides an upper bound on substrate-level resolution; the smooth-manifold limit emerges as $\xi \rightarrow 0$ in the appropriate scaling regime. This is closer to a falsification path (a sufficiently fine probe should reveal sub- ξ discreteness if VERSF is right) than a derivation, but it converts CGH_{smooth} from "axiomatically smooth" to "smooth as a tested continuum-limit prediction."

The present paper reduces CGH.0 but does not eliminate it. Dimension count and Lorentzian causal structure are derived under stated conditions. Smooth differentiable manifold structure remains a continuum-limit hypothesis, with a stated reduction direction. The remaining CGH burden is therefore narrowed from "assume 3+1 Lorentzian smooth manifold" to "assume a smooth continuum limit of the already-derived 3+1 Lorentzian causal structure" — a substantially smaller assumption with a concrete falsification probe.

Status table

Dependency	Old status (in upstream VERSF programme)	New status after closure
FH	Imported broad hypothesis (full fold-density programme)	Reduced to FH_min: no primitive substrate labels beyond commitment density
OC	Imported formalism (informal)	Reduced to OC-min: finite local protocols on committed records, with Lemma OC-1 closing pre-commitment access
Lorentz	Assumed invariant speed (CGH.1)	Regime characterisation + empirical branch selection; falsification = regime (a)
Spatial n = 3	Upstream result (assumed)	Imported as compact spatial admissibility result with stated argument structure
CGH.0	Assume 3+1 Lorentzian smooth manifold	Decomposed: metric structure derived; topology imported; smoothness reduced to continuum-limit hypothesis with $\xi \rightarrow 0$ reduction direction

What this paper does not derive

Three items lie outside the scope of this reduction and are flagged for completeness:

- **Field equations / dynamics.** The paper establishes the *kinematic envelope* within which VERSF dynamics must live. It does not derive Einstein's field equations or VERSF's dynamical equations of motion; those depend on additional capacity-loading and curvature-response inputs developed in the gravity sector of the programme.
- **Lorentzian signature.** Theorem 3 derives the Lorentz group from the existence of an invariant speed; it does not derive the (1,3) signature *a priori*. The signature is fixed by Theorem 1 (one timelike direction) plus D4 (three spatial directions).
- **Quantum sector.** The reduction is purely classical-kinematic. Quantum content (commitment statistics, the Born rule, etc.) is treated in the quantum-sector companion papers and does not appear here.

9. Conclusion

The standard postulates of relativistic kinematics — a single time dimension, Lorentz invariance, and the universality of free fall — are not independent. Each is the unique survivor of the A0 gauge sieve applied to a candidate physical structure that contains an operationally inaccessible component. Removing those components, A0 leaves behind exactly the relativistic kinematic content that the standard postulates assert directly.

The reduction is honest about its conditional structure. The Lorentz claim is a regime characterisation plus empirical selection. The WEP claim is conditional on FH_min, an upstream result still being formalised. The temporal-sector claim is unconditional under A0 + A2 + OC_min, with a precise statement of the bespoke quantity (temporal rank) that equals 1 — not the Dushnik–Miller order dimension. The remaining manifold burden is reduced from "assume 3+1 Lorentzian spacetime" to "assume a smooth continuum limit of the already-derived 3+1 Lorentzian causal structure". Every residual dependency is explicitly bounded in §8.

The unifying slogan:

The physical content of a structure is what survives every admissible observer-protocol. Whatever does not survive is not part of the world.

Time exists because facts exist; Lorentz invariance holds because no admissible protocol finds a substrate frame in the regime that committed records select; matter falls universally because gravity, as a substrate-level coupling, cannot distinguish what the substrate itself does not distinguish. The three are one reduction.

References

Foundational results invoked

- Bennett, C. H. (1982). The thermodynamics of computation — a review. *International Journal of Theoretical Physics* **21**(12), 905–940.
- Dushnik, B. and Miller, E. W. (1941). Partially ordered sets. *American Journal of Mathematics* **63**(3), 600–610. [Cited for contrast — order dimension is not the temporal rank used here.]
- Einstein, A. (1905). Zur Elektrodynamik bewegter Körper. *Annalen der Physik* **322**(10), 891–921.
- Landauer, R. (1961). Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development* **5**(3), 183–191.
- Szpilrajn, E. (1930). Sur l'extension de l'ordre partiel. *Fundamenta Mathematicae* **16**, 386–389. [Linear extension theorem; load-bearing in Theorem 1.]
- von Ignatowski, W. (1910). Einige allgemeine Bemerkungen zum Relativitätsprinzip. *Physikalische Zeitschrift* **11**, 972–976.

Related foundational programmes

- Bombelli, L., Lee, J., Meyer, D., and Sorkin, R. D. (1987). Space-time as a causal set. *Physical Review Letters* **59**(5), 521–524.
- Sorkin, R. D. (2005). Causal sets: Discrete gravity. In *Lectures on Quantum Gravity* (A. Gomberoff and D. Marolf, eds.), Springer, pp. 305–327. Proceedings of the 2002 Valdivia Summer School. arXiv:gr-qc/0309009.

Experimental tests of Lorentz invariance and the equivalence principle

- Mattingly, D. (2005). Modern tests of Lorentz invariance. *Living Reviews in Relativity* **8**, 5.
- Kostelecký, V. A. and Russell, N. (2011). Data tables for Lorentz and CPT violation. *Reviews of Modern Physics* **83**(1), 11–31. Annually-updated tables: arXiv:0801.0287 (latest version cited for current bounds).
- Touboul, P. et al. (MICROSCOPE collaboration) (2017). MICROSCOPE Mission: First Results of a Space Test of the Equivalence Principle. *Physical Review Letters* **119**(23), 231101.
- Touboul, P. et al. (MICROSCOPE collaboration) (2022). MICROSCOPE Mission: Final Results of the Test of the Equivalence Principle. *Physical Review Letters* **129**(12), 121102.
- Wagner, T. A., Schlamminger, S., Gundlach, J. H., and Adelberger, E. G. (2012). Torsion-balance tests of the weak equivalence principle. *Classical and Quantum Gravity* **29**(18), 184002.
- Will, C. M. (2014). The Confrontation between General Relativity and Experiment. *Living Reviews in Relativity* **17**, 4.

VERSF companion papers

- Taylor, K. *The fold-density derivation of bound-matter representation in VERSF*. VERSF Theoretical Physics Programme, AIDA Institute. [Source of the full Fold Hypothesis; FH_min used here is its minimal restriction.]
- Taylor, K. *A No-Go Theorem for Non-Simplicial Relational Substrates: Establishing $K = 7$* . VERSF Theoretical Physics Programme, AIDA Institute. [Structural input to FH_min.]
- Taylor, K. *The Origin of the Closure Scale in the VERSF Framework*. VERSF Theoretical Physics Programme, AIDA Institute. [Identifies the c of Theorem 2 with $\xi \times$ tick-rate; provides the $\xi \rightarrow 0$ reduction direction for CGH_smooth in D5.]
- Taylor, K. *The VERSF Uniqueness Theorem: The Fold as the Necessary Structural Core of a Minimal Fact-Producing Theory*. VERSF Theoretical Physics Programme, AIDA Institute.
- VERSF spatial-sector reductions (in preparation): the $n_{\text{space}} = 3$ result imported in D4 is developed in the upstream spatial-sector papers of the programme, drawing on BCB conservation, finite-capacity stability, and bound-state regularity arguments.