VERSF as the Origin of Special and General Relativity

A Constraint-Based Derivation from Record Dynamics and Finite Distinguishability

Keith Taylor — VERSF Theoretical Physics Programme

General-Reader Summary

Einstein's theories of special and general relativity tell us how space, time, and gravity work, but they begin by *assuming* certain things — that the speed of light is the same for everyone, that spacetime has four dimensions (three space, one time), that all matter falls the same way under gravity. These are postulates: the theories work brilliantly *if* you accept them, but the theories themselves don't explain *why* the universe should be that way.

This paper takes a step back and asks: what would have to be true for these assumptions to be *consequences* rather than starting points?

The answer the paper develops is that they emerge from much simpler, prior conditions about what it takes for *facts* to exist at all. Imagine the universe at its deepest level not as space and time but as a substrate that produces *committed records* — distinguishable, irreversible events that observers can compare. From four basic conditions on this substrate — that distinguishable states are finite, that records are irreversible, that any region holds a bounded amount of structure, and that influences spread locally — the paper shows the following:

- **The speed limit emerges naturally.** Because records can only spread so far in each tick of substrate time, there's a maximum speed at which information can move. Combined with the requirement that observers can't tell whose "rest frame" is more fundamental, this maximum speed becomes the universal invariant of special relativity.
- **Three dimensions of space come out uniquely.** In any other number of dimensions, atoms would be unstable, planetary orbits would be unstable, or gravity would behave so badly that bound structures couldn't exist. Three is the only dimension that allows the kind of stable matter we observe.
- **One dimension of time follows from how facts work.** A new foundational claim makes this precise: facts are committed records whose content and ordering are the same for all admissible observers. With zero time dimensions, you couldn't represent the irreversibility of records at all. With two or more time dimensions, observers could disagree about which event came first — but then those events wouldn't be facts at all (just observer-dependent commitments). Exactly one time dimension is the minimal structure that lets facts exist. The clean reading: *time exists because facts exist.*
- **Gravity behaves universally because matter is built from one fundamental ingredient.** The paper sketches an argument that all matter, regardless of type, is

composed of configurations of a single underlying substrate degree of freedom (the "fold"). Because that fold carries no labels distinguishing one type of matter from another beyond the fact of being committed, gravity must couple to all matter the same way — which is precisely the equivalence principle Einstein assumed.

- **General relativity emerges automatically.** Once you have the right speed limit, the right number of dimensions, and the right universal coupling, the standard mathematical arguments (which the paper uses but does not re-derive) deliver Einstein's equations.

The paper is honest about what it does and does not do. It does *not* re-derive relativity from scratch — Einstein's mathematical framework is used as-is. What is new is showing that the *inputs* Einstein had to assume are themselves consequences of deeper, simpler conditions about facts, records, and observers. Several of these reductions still rest on hypotheses that need further work in companion papers, and the paper flags every such dependency explicitly.

The paper also examines a candidate substrate architecture — a specific seven-vertex graph called the wheel W_7 — that realises the conditions abstractly required. This is exhibited as a working mechanism, not as the unique possibility; the paper is careful about scope here too.

The deep idea. Underlying the entire framework is a single principle the paper calls A0 (Observer-Invariant Distinguishability):

Only observer-invariant distinctions are real. Anything two admissible observers can disagree about — whether a particular event is moving, what type of matter is involved in some interaction, whether one event came before another — is not a fundamental physical fact. It is a feature of how observers describe things, not of what is actually there.

A natural question is whether this principle is itself an assumption — and if so, why accept it? The paper's answer is that A0 is not an arbitrary stipulation but the minimal condition required for the *concept* of a physical fact to be well-defined. A theory that allowed observer-relative quantities to count as physical would have no observer-independent answer to the question "what are the facts of the matter?" Different observers would find different answers, all equally valid, and physics would collapse into description. A0 is the consistency requirement that prevents this collapse — what every well-defined physical theory must satisfy if it is to have invariant content at all. So A0 is less an extra postulate of this paper than a clarification of what physics requires.

Applied to three different candidate distinctions, this single principle generates the three results above:

- Applied to *substrate frames* (combined with the irreversibility of records), it forces Lorentz invariance: the speed limit must be observer-invariant because the substrate's own preferred frame can never be detected by observers built from records.
- Applied to *matter types* (combined with the fold-substrate's single-degree-of-freedom content), it forces universal gravitational coupling: gravity cannot distinguish between matter types because the substrate cannot.

- Applied to *event orderings* (combined with the irreversibility of records), it forces exactly one temporal dimension: anything observer-relative about temporal order is, by A0, not a fact at all, and the substrate's actual facts admit only one consistent global ordering.

What emerges is a programme: relativity is not a set of postulates to be accepted but a set of consequences flowing from a prior level of structure where the only primitives are distinguishability, irreversibility, capacity, and locality, governed by the master principle that *only observer-invariant distinctions are real.* Space, time, and gravity are what that principle forces the world to look like. The paper takes that programme one substantial step forward, with every conditional dependency named.

Table of Contents

Abstract

We present a unified derivation of Special and General Relativity within the Void Energy–Regulated Space Framework (VERSF), a programme in which spacetime structure and gravitational dynamics emerge from constraints on distinguishability, irreversible commitment, and finite local capacity. The framework is axiomatised as a master principle A0 (Observer-Invariant Distinguishability) together with four substrate axioms (A1–A4) and four standing assumptions: (CGH) coarse-graining, (EFT) effective-theory minimality, (MARK) Markovian reduction, and (OC) observer-comparability — the last of which is implemented through the observer-protocol framework imported from the law-closure paper. A0 states that physical content depends only on distinctions invariant under all admissible observer-protocols; non-invariant distinctions are gauge and do not constitute physical facts. A0 is positioned not as an arbitrary axiom but as the minimal consistency requirement for a notion of physical fact to be well-defined: without it, physics would depend on observer-relative descriptions rather than invariant content.

Conditional on these, we establish four central results. First, Special Relativity arises from finite substrate propagation together with frame-invariance of the emergent propagation speed, with VERSF supplying the invariant speed that the standard von Ignatowski construction takes as input. Second, Newtonian gravitation arises when capacity loading varies spatially: a universal scalar sourcing of a gravitational potential by bound-record density, combined with the 3D Laplacian's Green's function, yields $\nabla^2\Phi = 4\pi G \rho_m$, with the operator form selected by a classification argument (Lemma 1) and equivalently presented as the Euler–Lagrange equation of a minimal action (Proposition 2). Third, General Relativity emerges as the unique covariant completion via the standard Lovelock and Deser–Feynman arguments, with VERSF supplying the substrate inputs those arguments require rather than re-deriving them. Fourth, the spatial dimension n = 3 is derived model-independently from BCB conservation, classical and quantum stability, and long-range regularity — with the temporal dimension supplied by A2's commitment ordering, conditional on (CGH.0)'s Lorentzian signature.

Theorem 9 (admissibility) and the K = 7 architectural examination are the two halves of the dimensional story, and we make their relationship explicit: **Theorem 9 selects 3+1 model-independently; the K = 7 wheel architecture provides a minimal candidate mechanism by which a VERSF substrate can realise that selection.** The two strands are logically independent: Theorem 9 does not depend on the K = 7 construction, and the K = 7 construction does not strengthen the admissibility result beyond realising it. Together they answer "is 3+1 forced?" and "can the substrate produce it?"

The K = 7 architectural results are scoped honestly. The substantive content of the $C_6 \rightarrow W_7$ transition is **topological-operational**: the hub introduces a topologically distinguished reference vertex (Lemma 2'), enabling radial-reference observer-protocols not realisable on a vertex-transitive boundary, which extends the admissible observable algebra. The corresponding **Revised Commitment Principle** is

Commitment = saturation of the admissible observable algebra by the constraint operator.

This supersedes earlier formulations framing commitment as elimination of physical null modes, which were inconsistent with what the connectivity lemmas actually establish.

Lemma 5's antipodal reduction is *not* C_6-specific: it generalises to any 6-vertex graph with an order-2 automorphism producing three antipodal orbits, including the octahedron and the prism. What distinguishes C_6 from these alternatives is **boundary edge-count minimality** (C_6: 6 edges; prism: 9; $K_{3,3}$: 9; octahedron: 12), supplemented by upstream criteria — fold-interface compatibility and primitive-eigenvalue normalisation — that operate at the upstream combinatorial level of the K = 7 No-Go Theorem and are not re-derived here. Every result labelled "derived within the C_6-boundary class" in §9 inherits conditionality on these upstream criteria; we make this dependency explicit.

A final section (§10) sketches three reductions of remaining standing assumptions to sharper substrate-level admissibility conditions, all unified as applications of the master principle A0. Lorentz invariance reduces to *no substrate-frame detection* (Theorem A: A0 + A1 + A2 + (OC) ⟹ NSF ⟹ Lorentz). WEP reduces to *no primitive gravitational charges* (Theorem B: A0 + fold-density-paper substrate content ⟹ NPC ⟹ universal coupling). The temporal dimension follows from the chain *facts exist ⟹ global comparability ⟹ one time dimension* (Theorem C: A0 + (OC) + A2 ⟹ Facts exist ⟹ GC ⟹ +1; the comparability condition itself is sourced from the §2.1 Proposition *Facthood = Invariant Comparability*, which is A0 applied to commitment events). Theorem D combines the spatial admissibility of Theorem 9 with Theorem C to fix the dimension count D = 3+1. The unifying reading: **(NSF), (NPC), and (GC) are not three independent admissibility conditions but three applications of A0 to substrate frames, record-type labels, and event orderings respectively.** These reductions do not yet remove the standing assumptions but localise each to a sharper residual hypothesis with concrete falsification paths.

---

1. Introduction

Special Relativity and General Relativity are extraordinarily successful theories that nonetheless rest on postulates the theories themselves cannot explain. Why an invariant signal speed? Why a four-dimensional Lorentzian manifold? Why the dimensionality we observe?

The VERSF programme begins several layers beneath the usual starting point of theoretical physics. Rather than postulate an invariant speed, a curved spacetime, or a specific dimensionality, VERSF asks what structure must be in place for physical facts to exist at all. The answer comes in four pre-geometric axioms about distinguishability, commitment, capacity, and locality, none of which mentions space, time, dimension, or causal structure.

1.1 Three layers of structure

Layer 1 — the substrate. A0 + A1–A4 live here. A0 (Observer-Invariant Distinguishability) is the master principle: physical content depends only on distinctions invariant under all admissible observer-protocols. A1–A4 are the substrate axioms (finite distinguishability, irreversible commitment, finite capacity, local coupling). They are pre-geometric: no manifold, no dimension, no metric. Layer 1 is developed in companion papers (the K = 7 No-Go Theorem, the commitment-barrier paper, the fold-density paper); the present paper treats it as foundation.

Layer 2 — emergent geometry. Where dimension and manifold structure appear. We address Layer 2 in two complementary ways that together form a single coherent argument:

- The *model-independent* strand (Theorem 9) shows that any substrate of the relevant type admits only 3+1 dimensions. This answers the question "*is* 3+1 forced?"
- The *substrate-specific* strand examines a candidate VERSF closure cell (the C_6-wheel) as a *realisation* of this admissibility. This answers the question "*can* a VERSF substrate produce it?"

These are not parallel claims; they are the existence and realisation halves of a single dimensional story. **Theorem 9 selects 3+1; the K = 7 wheel architecture provides a minimal candidate mechanism by which a VERSF substrate can realise that selection.** The two strands are logically independent: Theorem 9 does not depend on the K = 7 construction, and the K = 7 construction does not strengthen the admissibility result beyond realising it. Without Theorem 9, the K = 7 construction would be only suggestive — three dimensions emerging by combinatorial accident. Without the K = 7 construction, Theorem 9 would establish only that 3+1 is admissible, without exhibiting any substrate that produces it. Together, the two strands close the loop: a substrate with the right pre-geometric structure exists, and the only dimension it can produce is the one observed.

Layer 3 — effective physics. Once Layer 2 supplies a manifold of the right dimensionality, the physics of records propagating on it is straightforwardly SR + Newtonian gravity + GR. The derivations here use standard arguments (Lovelock, Deser–Feynman) wherever standard arguments suffice; VERSF's contribution at Layer 3 is to supply these standard arguments with the inputs they require from Layer 1.

1.2 Three corrections to the framing of earlier VERSF work

Three corrections to earlier framings — both in earlier drafts of this paper and in the wider VERSF programme literature — warrant prominence here, because they shape how the substrate-specific results in §9 should be read:

(i) Commitment is saturation of the admissible observable algebra, not elimination of null modes. Earlier formulations framed commitment as elimination of physical null modes under hub coupling. This is inconsistent with what the connectivity lemmas (2 and 3) actually establish — both null modes are gauge-invariant uniform shifts in the standard sense — and is retracted. The Revised Commitment Principle (§9.4.1) supersedes it: a cell is committed when its constraint operator saturates the admissible observable algebra, i.e., when rank(B_G) = dim($\mathcal{O}$(G)).

(ii) The $C_6 \to W_7$ transition is topological-operational, not spectral. The substantive content is the hub's introduction of a topologically distinguished reference vertex (Lemma 2'), enabling new observer-protocols (Definition 6.1) that extend the admissible observable algebra. The rank/nullity content of Lemmas 2 and 3 is supportive but not the primary mechanism.

(iii) The K = 7 architecture realises 3+1 admissibility but does not derive 3+1 in isolation. Dimensional admissibility is selected model-independently by Theorem 9 plus (CGH.2) isotropy. The K = 7 construction is the architectural mechanism by which this admissibility is realised in the VERSF substrate, conditional on the upstream selection of C_6 as the canonical boundary among antipodal-paired 6-vertex graphs (octahedron, prism, $K_{3,3}$-with-pairing).

Footnote (companion-paper propagation). The retractions in (i) and (ii) propagate to claims in companion papers within the VERSF programme that have used the "spectral signature of commitment" or "physical-to-gauge reclassification" framings. These references should be updated to align with the Revised Commitment Principle and topological-operational characterisation introduced here. Specifically: the commitment-barrier paper and any portion of the fold-density paper invoking the retracted framings should be revised. The K = 7 No-Go Theorem rests on combinatorial criteria (edge-count minimality, fold-interface compatibility, primitive-eigenvalue normalisation; see §9.10) rather than on the commitment-mechanism framing being retracted; we believe its arguments survive the corrections of (i) and (ii) intact, but a full audit of the No-Go Theorem against the Revised Commitment Principle has not been completed and is itself an open programme item. Until that audit is done, claims in this paper that depend on the No-Go Theorem (notably the §9.10 selection of C_6) carry the implicit condition that the No-Go Theorem's derivation is not load-bearing on the retracted framings.

1.3 Nature of the contribution

A fair reader might ask what is novel here. Two kinds of contribution should be distinguished.

Novel original derivation:

- **Lemma 4' (Radial Observable Independence)** — a new linear-algebra argument establishing that the radial observable on W_7 is genuinely independent of boundary edge-differences, closing the algebraic gap underlying $\dim(\mathcal{O}(W_7)) = 6$.
- **The Revised Commitment Principle (§9.4.1)** — *commitment = saturation of the admissible observable algebra* — replacing the earlier null-mode framing that was inconsistent with what the connectivity lemmas establish.
- **The topological-operational mechanism** of the $C_6 \to W_7$ transition: the hub's introduction of a topologically distinguished reference vertex (Lemma 2') enabling radial-reference observer-protocols, with the resulting observable-algebra extension as the primary content (Theorem 13B).
- **The honest scope-mapping of §9.10**, separating criterion (A) edge-count minimality (exhibited here) from criteria (B) and (C) (upstream-asserted), with the conditionality of all §9 results made mechanical.

Novel framing of standard results:

- **Theorem 9 against the Tangherlini–Tegmark line.** The classical–quantum stability case for $n = 3$ has been observed before. What is new is to embed these arguments within a substrate programme where the inputs they take as ambient (Gauss's law, the Schrödinger equation) are themselves substrate-derived: BCB conservation comes from A1, A2 and the record-current structure rather than from classical electromagnetism. See §9.2's added-value remark for a fuller account.
- **Theorem 7 against the Lovelock + Deser–Feynman line.** GR's covariant completion is standard. What is new is the argument that the inputs Lovelock and Deser–Feynman take as premises — Lorentz-invariant propagation, universal scalar sourcing, $D = 4$, universal stress-energy coupling — are themselves consequences of A1–A4 plus the standing hypotheses, rather than postulates about geometry.

The paper does not claim to derive all of physics. It claims to derive SR + GR + 3+1 admissibility from pre-geometric axioms, under stated conditions, with every dependency made explicit, and to exhibit a candidate substrate that realises the admissibility.

1.4 Organisation

§2 collects axioms and standing assumptions. §3 introduces record dynamics. §§4–7 develop SR, Newtonian gravity, GR, and the unified picture. §8 takes stock. §9 addresses the dimensional programme: model-independent admissibility (§9.2) and the substrate-specific examination (§§9.3–9.11). §10 sketches three reductions of the remaining standing assumptions to sharper substrate-level admissibility conditions — these are forward-looking strengthenings rather than core derivations, labelled Theorems A, B, C, D outside the main Theorem 1–13 sequence. §11 concludes.

---

2. Axioms and Standing Assumptions

2.1 The VERSF axioms

A0 (Observer-Invariant Distinguishability) — Master Principle. Let $\mathcal{O}$ be the admissible observable algebra (Definition 6.0) and let states be $\mathbf{x} \in \mathbb{R}^{|V|}$ over the relevant cell. Define the equivalence relation

$$\mathbf{x} \sim \mathbf{y} \Leftrightarrow \forall O \in \mathcal{O},\ O(\mathbf{x}) = O(\mathbf{y}).$$

Then physical content depends only on equivalence classes under ∼. Equivalently:

Physics depends only on distinctions invariant under all admissible observer-protocols; non-invariant distinctions are gauge and do not constitute physical facts.

A0 is the master principle of the programme. It states a single criterion: the only distinctions that count as physical are those preserved by every admissible observer-protocol. Distinctions that some admissible protocol cannot resolve, or that admissible protocols disagree about, are gauge — they may exist in some lower-level description of the substrate, but they do not constitute physical facts at the level where physics is to be done.

A0 as consistency requirement, not arbitrary axiom. A0 is not introduced as an arbitrary axiom but as the minimal condition required for a notion of physical fact to be well-defined: if non-invariant distinctions were allowed, physics would depend on observer-relative descriptions rather than invariant content, and there would be no observer-independent answer to "what are the facts of the matter?" A theory that classifies observer-relative quantities as physical loses the property that distinguishes physics from description; A0 is the consistency condition that prevents this collapse. The principle is therefore not a strong claim about how the world is but the minimal commitment without which the concept of a physical fact has no observer-invariant meaning. A reader inclined to ask "why accept A0?" should ask instead "what would it mean to do physics without it?" — and finding no satisfactory answer, will see A0 as the requirement implicit in any well-defined physical theory rather than as an extra postulate of this one.

Operational characterisation of gauge. A0's "non-invariant distinctions are gauge" clause has a precise operational meaning that ties to the algebraic machinery of §9.4: *gauge distinctions are precisely those that vary under admissible observer-protocol transformations while leaving all observable outputs invariant.* In the Definition 6 framework, this is the relative gauge/physical classification: a vector $v \in \ker(B_G)$ is gauge relative to $\mathcal{O}(G)$ iff every admissible observer-protocol on G takes the same value on $\mathbf{x}$ and $\mathbf{x} + v$ (Lemma 4, Kernel–Protocol Correspondence). A0 is therefore not free-standing; it is the principle that selects which equivalence classes under this relation count as physical content. The kernel/quotient machinery already in the paper — Lemma 4, the Saturation criterion ($\mathrm{rank}(B_G) = \dim \mathcal{O}(G)$), the Gauge Quotient Lemma — is the operational implementation of A0 at the cell level.

A0's role in the architecture. A0 is foundational rather than derived. It precedes A1–A4 not in a temporal sense but in a logical one: A0 specifies *which substrate-level distinctions count as physical content*, while A1–A4 specify *what the substrate is and how it behaves*. A0 + A1–A4 together constitute the programme's full set of premises, with A0 supplying the criterion of physicality and A1–A4 supplying the substrate dynamics.

A0 has substantive consequences. The reductions of §10 — Theorem A's (NSF), Theorem B's (NPC), Theorem C's (GC) — are not three independent admissibility conditions but three applications of A0 to specific candidate distinctions:

- **NSF (Theorem A).** A0 forbids physical distinctions that admissible observer-protocols cannot detect; the substrate update frame, by A1 + A2, is precisely such an undetectable distinction; therefore frame-relative quantities are gauge, and the effective transformation group preserves the causal cone (Lorentz invariance).
- **NPC (Theorem B).** A0 forbids physical distinctions between record types that admissible observer-protocols cannot independently resolve; if record types are effective-field labels of fold configurations rather than primitive substrate charges (fold-density paper), then differences between the κ_a coupling coefficients would be A0-forbidden; therefore $\kappa_a = \kappa$ (universal coupling).

- **GC (Theorem C).** A0 forbids physical distinctions in the temporal order of commitment events that admissible observers disagree about; therefore the substrate's *facts* (Facthood = Invariant Comparability, below) are globally comparable; therefore the minimal coarse-grained representation has exactly one timelike dimension.

The unifying reading: **NSF, NPC, and GC are not separate hypotheses but three consequences of A0 applied to substrate frames, record types, and event orderings respectively.** Each consequence requires additional substrate input (A1+A2 for NSF; fold-density for NPC; A2 + Facthood for GC), but the operative principle in each case is A0. We make this explicit in §10's reduction statements.

A0 and the Facthood Proposition. The Facthood = Invariant Comparability Proposition (below) is a special case of A0 applied to commitment events: where A0 says only invariant distinctions are physical, the Facthood Proposition says facts are committed records invariant under all admissible observer-protocol comparisons. The Proposition is preserved as a separate statement because it is the formulation actually used in §10.3 and because it operationalises A0 specifically for commitment events; it is not redundant with A0 but its commitment-event-specific corollary.

A1 (Finite Distinguishability). In any finite region of the substrate, only finitely many physical states can be operationally distinguished.

A1 is the foundation. The world is not continuous at its deepest level; there is a last layer at which states can be told apart, and it is finite.

A2 (Irreversible Commitment). The substrate supports stable records produced by irreversible transitions; erasure of a record requires nonzero thermodynamic cost.

A2 is the programme's thermodynamic axiom. Its content is Landauer-like. We will later encounter a separate, structural notion of "commitment" (Definition 7) defined through observable-algebra saturation; the relation between A2's thermodynamic commitment and Definition 7's structural commitment is the subject of an open conjecture in §9.5.

Proposition (Facthood = Invariant Comparability). *A physical fact is a committed record whose content and ordering are invariant under all admissible observer-protocol comparisons.*

Objects failing this criterion are observer-dependent commitments, not physical facts.

This proposition is the commitment-event-specific corollary of A0: where A0 specifies that only observer-invariant distinctions are physical, the Facthood Proposition specifies that *facts* — the substrate content most directly relevant to physics — are exactly the commitment events satisfying A0's invariance criterion. The proposition does not introduce new content beyond A0; it operationalises A0 for the specific category of substrate content (commitment events) used in Theorem C.

The proposition is used substantively in §10.3 (Theorem C). The conceptual upgrade is significant: the existence of a temporal dimension at all becomes a consequence of the existence of facts. Where v25's framing read "(OC) + A2 ⟹ (GC) as a derived implication," v29's framing reads:

Facts exist ⟹ (GC) holds ⟹ exactly one temporal dimension.

This connects ontology (what counts as a fact), thermodynamics (A2's irreversibility), observer theory ((OC)), and geometry (the temporal dimension) into a single argument chain. The clean form: *time exists because facts exist.*

A3 (Finite Capacity). Every finite region has a bounded capacity to sustain distinguishable committed structure. Writing $\Sigma(x)$ for the local loading and Σ_c for the threshold, the dimensionless loading is $u(x) := \Sigma(x)/\Sigma_c \in [0, 1]$.

A3 prevents unbounded record formation and gives rise, through the proper-time response function, to gravitational phenomena.

A4 (Local Coupling). Each elementary substrate update affects only a bounded neighbourhood.

A4 supplies a notion of "nearness" prior to any notion of distance; combined with A1 it ensures record-relevant influence cannot propagate arbitrarily fast.

A1–A4 describe a substrate of finite distinguishability, irreversible record formation, bounded capacity, and locality. None mentions dimension, manifold, metric, or spacetime.

2.2 Standing assumptions

(CGH.0) Emergent manifold. Coarse-graining produces a smooth Lorentzian manifold $\mathscr{M}$. Dimension not assumed; derived in §9.2.

(CGH) Coarse-graining of dynamics. (CGH.1) c is frame-invariant on $\mathscr{M}$; (CGH.2) field representatives are smooth, with the IR theory homogeneous and isotropic; (CGH.3) the substrate's preferred frame is operationally inaccessible; the IR theory is diffeomorphism-invariant.

(EFT) Effective Field Theory minimality. Local, second-order, standard EFT power-counting.

(MARK) Markovian reduction. Coarse-grained record dynamics is Markovian.

(OC) Observer-comparability via observer-protocols. Physical variables correspond to stable, observer-comparable distinctions. The operational implementation, used in §9 to interpret substrate architectures, is the observer-protocol framework of the companion law-closure paper. Observer-protocols are computational procedures executable on graph cells using physically realisable measurement primitives.

Important scope statement. The observer-protocol framework — including the specification of "physically realisable measurement primitives" — is *axiomatically specified by the law-closure paper, prior to and independently of any notion of commitment*. The primitives are not defined via commitment events, A2-thermodynamic commitment, or Definition 7's structural commitment; they are stipulated externally. This is the fence-post that prevents Definition 6.1's admissibility from sliding back into circularity. Every claim in this paper that depends on observer-protocols inherits this conditionality on the law-closure paper. A substrate-first derivation of (OC) and its primitives from A1–A4 is an open programme target.

Forward reference to §10. Section 10 sketches three reductions of standing assumptions — (CGH.1) frame-invariance via Theorem A, WEP via Theorem B, the temporal-dimension argument via Theorem C — to sharper substrate-level admissibility conditions. These reductions are forward-looking strengthenings that narrow what (CGH), WEP, and Theorem 9 Step 5 must hypothesise; they do not remove the standing assumptions but make their residual content sharper. A reader of §2.2 should know §10 exists; the present sections work entirely within the original standing-assumption framing for clarity, with the strengthenings as a consolidating final section.

2.3 Why (CGH) is a credible assumption

(CGH) carries a significant share of the paper's inferential load. It is not arbitrary; it has well-studied physical precedent in BEC analogue gravity (Unruh 1981; Barceló–Liberati–Visser 2011), where phonon excitations propagate on an emergent Lorentzian metric with a universal limiting speed (the local sound speed) frame-invariant for all phonon-scale observers, even though the condensate's microscopic rest frame is operationally inaccessible. (CGH.1) and (CGH.3) assert an analogous mechanism for the VERSF substrate. A substrate-first proof of (CGH) from $K = 7$ simplicial structure is a separate open programme target. This paper derives consequences given (CGH), not (CGH) itself.

Minimal weakening of CGH for §§4–7. It is worth noting that the relativity derivations of §§4–7 do not require the full strength of (CGH.0)'s smooth differentiable manifold structure at leading order. They require only: Lorentz-invariant causal cones (used in Theorem 2 and §6); smooth coarse-grained fields (used in Theorems 3, 5 and the Taylor expansion of Proposition 1); and isotropy plus diffeomorphism invariance at the IR scale (used in Theorem 7). Higher-order or strong-field results would require the full manifold structure, but the weak-field regime can in principle be developed on a less restrictive emergent geometry. This minimal-weakening observation is included for completeness; the present paper assumes (CGH.0) in full for clarity rather than working at the boundary of what can be derived from less.

---

3. Record Dynamics and Conservation Structure

A *record* is a stable, irreversible, independently comparable physical distinction (consequences of A1–A4). An influence is **record-relevant** if it can alter the future distribution of committed records.

The record current $C^\mu := (\rho_c, \mathbf{J}_c)$ on $\mathcal{M}$ satisfies $\partial_\mu C^\mu = s_c$, with s_c instantaneous-only under (MARK).

---

4. Special Relativity

§§4–7 work throughout on a 3+1-dimensional Lorentzian manifold $\mathcal{M}$ whose dimensional admissibility is established model-independently in Theorem 9 (§9.2). Results in these sections are conditional on that dimension selection.

Special Relativity is the kinematics of finite, invariant record propagation. Theorem 1 supplies finite speed (from A1, A4); (CGH.1) supplies invariance; the Lorentz group follows.

Theorem 1 (Finite Substrate Propagation Speed)

Premises: A1, A4.

A1 + A4 give bounded neighbourhood r_0 per update; $c := a\, r_0 / \Delta t$ is finite. ■

Theorem 2 (Lorentzian Kinematics)

Premises: A1, A4, (CGH).

(CGH.3) gives group structure, (CGH.2) gives linearity, (CGH.1) gives causal-boundary preservation, (CGH.2) gives isotropy. The von Ignatowski argument fixes γ; the result is the Lorentz transformation. ■

Corollary 1 (Minkowski Structure).

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

---

5. Newtonian Gravity

Gravitation is the response of a finite-capacity substrate to non-uniform commitment loading. The argument arc: bound records (Definition 4) and capacity loading (Definition 5) feed into bound-record continuity (Theorem 3) and the proper-time response (Theorem 4). The Taylor expansion (Proposition 1) defines the gravitational potential. Lemma 1 selects the Laplacian as the unique admissible scalar field operator under the symmetries; combined with continuity, this gives Poisson's equation (Theorem 5), with Proposition 2 presenting the same equation as the Euler–Lagrange equation of a minimal action — a Lagrangian restatement of the operator-classification result rather than a logically independent derivation. The metric component identification $g_{00} = R(u)^2$ (Proposition 3) bridges to GR. Geodesic motion (Theorem 6) recovers Newton's second law in the slow-motion limit.

Definitions and Theorems

Definition 4. $\rho\_m := \mu\, \rho\_bound$.

Definition 5. $\Sigma = \kappa_1\, \rho\_bound$ (5a) or $\int K(|x-y|)\, \rho\_bound\, d^3y$ (5b); $K(r) \propto 1/r$ is the Newtonian case.

Theorem 3. $\partial\_t\, \rho\_bound + \nabla\cdot\mathbf{J}\_bound = S\_commit$ (consequence of A2, smoothness, MARK).

Theorem 4. $d\tau/dt = R(u(x))$, $R(0) = 1$, $R'(u) < 0$.

Proposition 1. $R(u) = 1 - \alpha u + \beta u^2 + \mathcal{O}(u^3)$; $\Phi/c^2 := R - 1$; at first order $\Phi = -\alpha c^2 u$; $d\tau = (1 + \Phi/c^2)dt + \mathcal{O}(\Phi^2/c^4)$.

Physical motivation for the constitutive relation. Higher-loading regions are proper-time sinks; bound records drift toward them because commitment per proper-time tick is more efficient there. $\mathbf{J}\_bound = -\chi\, \nabla\Phi$ with $\chi > 0$.

Lemma 1 (Poisson Uniqueness in d = 3). Under locality, isotropy, linearity, staticity, ellipticity with vanishing Green's function at infinity, and ≤ 2 derivatives, $\nabla^2\Phi = C\rho$ uniquely. *Proof: operator span $\{1, \nabla^2\}$; identity excluded by ellipticity; Helmholtz excluded by EFT; biharmonic excluded by order.*

Remark on the d ≥ 3 dependence. The hypothesis "ellipticity with vanishing Green's function at infinity" silently encodes $d \geq 3$: in $d = 2$ the Laplacian's fundamental solution is logarithmic and does not vanish at infinity; in $d = 1$ it grows linearly. Lemma 1 is therefore a uniqueness result *given* that $d = 3$ (or more broadly $d \geq 3$ with the regularity hypothesis), not a derivation of $d = 3$. The latter is the role of Theorem 9. Lemma 1 cannot be invoked to derive $d = 3$ from itself.

Theorem 5 (Poisson Equation). $\nabla^2\Phi = 4\pi G\, \rho\_m$. *Proof: static continuity $\Longrightarrow -\chi\nabla^2\Phi = S$; admissibility gives $S = -\alpha\_S\, \rho\_bound$; combining, $\nabla^2\Phi = (\alpha\_S/\chi\mu)\rho\_m$, with $4\pi G := \alpha\_S/(\chi\mu)$.*

Remark (WEP). Universal scalar sourcing — that every bound record contributes identically to capacity strain — is the weak equivalence principle in substrate language. VERSF names rather than derives this. A3 supplies the scalar-functional form $\Sigma = F[\rho\_bound]$ but does not by itself entail that the functional is universal across bound-record types: in principle, F could weight different record types non-uniformly (electromagnetic vs nuclear vs mechanical, in the eventual matter-sector decomposition), corresponding to a violation of WEP. The assumption that F is universal — that every bound record contributes identically to the loading regardless of its internal structure — is an additional input beyond A3's scalar-loading structure. A substrate-first derivation of universal coupling, from a structural argument forcing F to be record-type-blind, would upgrade Theorem 5 from conditional to unconditional.

Proposition 2. $S[\Phi] = \int[(1/8\pi G)\cdot\frac{1}{2}|\nabla\Phi|^2 + \Phi\rho\_m]d^3x$; $\delta S/\delta\Phi = 0 \Longrightarrow$ Poisson.

Corollary 2. $\Phi(r) = -GM/r$; $\mathbf{g}(r) = -(GM/r^2)\hat{\mathbf{r}}$.

Proposition 3. $g_{00}(x) = R(u(x))^2$. *Proof: $d\tau/dt = R \Longrightarrow d\tau^2 = R^2\, dt^2$; rest-frame line element $d\tau^2 = g_{00}\, dt^2$; hence $g_{00} = R^2$.*

Expanding, $g_{00} = 1 + 2\Phi/c^2 + \mathcal{O}(u^2)$. SR $\leftrightarrow$ $R \equiv 1$ (Minkowski); GR $\leftrightarrow$ $R(u)$ varies.

Consistency remark. The identification $g_{00} = R(u)^2$ is consistent with Lorentz invariance (Theorem 2) and preferred-frame inaccessibility (CGH.3): R varies as a function of the loading u(x), which is a coordinate-scalar built from substrate-density data, not as a function of any observer's velocity. A Lorentz boost transforms u(x) covariantly — through the boost of x — and produces the same R(u(x')) at the boosted point, with no preferred frame singled out by R's variation. Lorentz invariance applies to the kinematic structure on $\mathscr{M}$; R's spatial variation is dynamical content compatible with that structure rather than a violation of it. This is the standard interpretation: spatially varying gravitational fields curve spacetime without breaking the local Lorentz structure of the tangent space at each point.

Remark on falsifiable post-Newtonian content. The Taylor expansion of R(u) is the natural place where post-Newtonian deviations would enter VERSF. Substituting $\Phi/c^2 = R - 1$ into $g_{00} = R^2$ and expanding to second order in u:

$$g_{00} = 1 - 2\alpha u + (\alpha^2 + 2\beta)u^2 + \mathcal{O}(u^3) \qquad = 1 + 2(\Phi/c^2) + (1 + 2\beta/\alpha^2)(\Phi/c^2)^2 + \mathcal{O}((\Phi/c^2)^3).$$

In the standard parametrised post-Newtonian (PPN) form, $g_{00} = 1 + 2(\Phi/c^2) + 2\beta\_PPN(\Phi/c^2)^2 + \ldots$, where β_PPN = 1 in GR and is constrained empirically to within $\sim10^{-4}$ of unity by tests including Cassini's Shapiro delay and lunar laser ranging. Identification gives

$$\beta\_PPN = \tfrac{1}{2} + \beta\_VERSF/\alpha^2.$$

So VERSF predicts β_PPN = 1 (the GR value) precisely when $\beta\_VERSF = \alpha^2/2$. Any other relationship between the substrate coefficients α and β produces a deviation from GR in solar-system observables. This is a *conditional* falsifiable prediction: if a future substrate-first derivation fixes the ratio β/α^2 to a value other than 1/2, the discrepancy would be measurable; if the ratio is fixed at 1/2 by substrate considerations, β_PPN = 1 becomes a substrate-derived consequence rather than a phenomenological match. The PPN parameter γ (governing light deflection and Shapiro delay through g_ij) is determined in this paper by Theorem 7's spin-2 universality rather than by R(u)'s expansion, and is therefore predicted as γ = 1 unconditionally within the Lovelock + Deser–Feynman covariant completion of §6.

This connects the substrate-level expansion of R(u) directly to solar-system tests, providing one concrete falsifiable interface between VERSF substrate parameters and observed gravity beyond the Newtonian regime.

Theorem 6. Slow-motion limit: $d^2\mathbf{x}/dt^2 = -\nabla\Phi$. The factor of 2 in $g_{ij} = -(1 - 2\Phi/c^2)\delta_{ij}$ is fixed by Theorem 7's spin-2 universality.

---

6. General Relativity

Theorem 7 (Effective Covariant Completion)

Premises: Theorem 5, Theorem 6, Proposition 3, Corollary 1, D = 4 (Theorem 9), (CGH), (EFT).

$S_{grav} = (1/16\pi G) \int (R - 2\Lambda) \sqrt{|g|}\, d^4x$; $G_{\mu\nu} + \Lambda g_{\mu\nu} = (8\pi G/c^4) T_{\mu\nu}$.

Proof. *Step 1 (Lovelock in D = 4).* Diffeomorphism invariance + EFT second-order ⟹ $G_{\mu\nu} + \Lambda g_{\mu\nu}$ is the unique admissible tensor. *Step 2 (Deser–Feynman).* Lorentz covariance + massless spin-2 + universal coupling + nonlinear self-consistency bootstrap ⟹ Einstein–Hilbert. Newtonian matching fixes G. ■

Remark on the Deser–Feynman bootstrap. The bootstrap argument has known subtleties that we do not resolve here. The bootstrap establishes the Einstein–Hilbert structure to all orders in $h_{\mu\nu}$ assuming on-shell consistency and the absence of further-derivative completions; off-shell ambiguities and higher-derivative extensions modulo field redefinitions are points of continuing technical discussion in the literature. The (EFT) hypothesis used in Step 1 — at most second-order in metric derivatives — addresses the higher-derivative question by stipulation; the off-shell ambiguities are typically argued to be field-redefinition artefacts that do not affect physical predictions in the IR regime relevant to this paper. We invoke the bootstrap in its standard form and inherit its standard subtleties; readers seeking rigorous treatment of the bootstrap's uniqueness should consult the technical GR literature.

Remark (nature of contribution). VERSF does not re-derive Lovelock or Deser–Feynman. The VERSF contribution is to supply substrate-level inputs they take as premises — Lorentz-invariant propagation (Theorems 1, 2), universal scalar sourcing (Theorem 5), D = 4 emergent dimensionality (Theorem 9), universal stress-energy coupling (matter sector). Standard GR derivations cannot pose "where do these inputs come from?" within their own scope; VERSF can.

Remark (Λ). Lovelock fixes the form of Λ but not its value; substrate-first origin is companion-paper work.

---

7. The Unified Picture

Theorem 8 (VERSF Relativity Correspondence)

Uniform record propagation ⟹ SR; non-uniform capacity ⟹ GR; unified via $g_{00} = R(u)^2$. ■

---

8. Discussion

8.1 Status of every substantive claim

(Single entry per claim; supporting/derivative results are noted in their parent row. The "C_6-class" annotation marks results conditional on the upstream selection criteria of §9.10.)

Result	Source	Conditional on	Status
A0 (Observer-Invariant Distinguishability) — Master Principle	**§2.1**	minimal consistency requirement for a well-defined notion of physical fact	**Foundational programme principle** (generates NSF, NPC, GC of §10; not an arbitrary axiom but the consistency condition without which physical facts have no observer-invariant meaning)
Smooth $\mathscr{M}$	(CGH.0)	—	Standing hypothesis
3+1 admissibility (model-independent)	**Theorem 9**	A1–A4, stability, regularity	**Derived**
Frame-invariance of c	(CGH.1)	—	Standing hypothesis (analogue precedent)
Smoothness/isotropy of IR	(CGH.2)	—	Standing hypothesis
Diffeomorphism invariance	(CGH.3)	—	Standing hypothesis
EFT minimality	(EFT)	—	Standing hypothesis
Markovian reduction	(MARK)	—	Standing hypothesis
Observer-comparability via observer-protocols	(OC)	law-closure paper (axiomatic primitives)	**Imported (structural assumption)**
Facthood = Invariant Comparability	**§2.1 Proposition**	A1 + A2 + (OC)	**Foundational programme content** (defines hard boundary: facts vs observer-dependent commitments)
Substrate speed c finite	Theorem 1	A1, A4	Derived
Lorentz transformations	Theorem 2	(CGH)	Derived
Bound-record continuity	Theorem 3	A2, (CGH.2), (MARK)	Derived
R(u) Taylor form	Proposition 1	smoothness	Derived

Result	Source	Conditional on	Status
Universal scalar sourcing (WEP)	Theorem 5 premise	—	**Named, not derived**
Poisson uniqueness in d = 3	Lemma 1	(EFT), d = 3	Derived
$\nabla^2\Phi = 4\pi G \rho_m$	Theorem 5	WEP, Lemma 1, (CGH)	Derived (conditional on WEP)
Variational form	Proposition 2	Theorem 5 premises	Derived
$g_{00} = R(u)^2$	Proposition 3	Theorem 4 + Lorentzian rest frame	Derived
Einstein–Hilbert action	Theorem 7	(CGH), (EFT), D = 4	Derived
Einstein field equations	Theorem 7	same	Derived
Revised Commitment Principle	**§9.4.1 (boxed)**	(OC) framework	**Stated as principle**
Saturation criterion	§9.4.5	rank(B_G) = dim($\mathcal{O}$(G))	Definition
Observable algebra structure (vector space)	**Definition 6.0**	linearity of protocols + observational equivalence	**Supplied**
Observer-protocol admissibility	Definition 6.1	law-closure paper primitives	Imported
Observable null mode (relational)	Definition 6	Definition 6.1	Supplied
Kernel–protocol correspondence	Lemma 4	Definitions 6, 6.1	Derived
Radial observable independence	**Lemma 4'**	direct linear-algebra	**Derived**
Antipodal reduction 6 → 3 (generic)	Lemma 5	(CGH.2)	Derived; not C_6-specific
Commitment (structural)	Definition 7	Revised Commitment Principle	Definition
Automorphism asymmetry of C_6 vs W_7	Lemma 2'	graph theory	Derived (foundational)
Connectivity of B_6, W_7	Lemmas 2, 3	—	Derived (connectivity-level only)
Observable-algebra extension under hub coupling — primary mechanism	**Theorem 13B**	Lemma 2', Lemma 4, Lemma 4', (OC), C_6-class	**Derived (primary result of §9, conditional on §9.10)**

Result	Source	Conditional on	Status
Topological-operational distinction (auxiliary)	Theorem 12	Lemma 2', Lemma 4, (OC), C_6-class	Derived (auxiliary, not required for Theorem 13B)
A2 ↔ Definition 7 thermodynamic bridge	**§9.5**	—	**Open conjecture**
K = 7 wheel realises 3+1 admissibility	Theorem 10	Theorem 9, C_6-class	Derived (conditional on §9.10)
K < 7 excluded within C_6-class	Theorem 11	C_6-boundary admissibility	Derived (conditional on §9.10)
Structural minimality (Theorem 13A)	C_6-class	Lemmas 5, 11	Derived (conditional on §9.10)
Operational/thermodynamic (Theorem 13C)	C_6-class	§9.5 conjecture	Conjecture (conditional on §9.10)
Edge-count minimality of C_6 among antipodal-paired 6-vertex graphs	§9.10(A)	direct count	Exhibited as preference
Fold-interface compatibility favouring C_6	§9.10(B)	upstream K=7 No-Go Theorem	Asserted upstream, partially described, not re-derived
Primitive-eigenvalue normalisation favouring C_6	§9.10(C)	upstream K=7 No-Go Theorem	Asserted upstream, parsimony-type argument, not re-derived
K ≥ 8 uniqueness in C_6-class	§9.11	bounded enumeration	Open
Theorem A: Lorentz from (NSF)	§10.1	(NSF) — itself sourced from A1 + A2 + (OC)	**Strong conditional theorem**
Theorem B: WEP from (NPC)	§10.2	(NPC) — itself sourced from upstream fold-density paper content	**Strong reduction**
Theorem C: Single time from (GC)	§10.3	(GC) — itself sourced from (OC) + A2; constitutive of facthood per §2.1 Proposition	**Plausible theorem requiring formalisation**
Theorem D: D = 3+1 dimension count	§10.4	Theorems 9 + C	**Derived combination** (dimension count only)
K = 7 as derivation of 3+1 in isolation	—	—	**Not claimed** (K = 7 realises Theorem 9)

9. The Dimensional Programme

A1–A4 contain no reference to dimension; the emergence of 3+1 must be derived. We address this in two strands that together form a single coherent argument.

9.1 Overview

Model-independent strand (§9.2). Theorem 9: $n = 3$ uniquely admissible. This is the *existence* result: it shows that no other dimension is compatible with A1–A4 plus standard stability and regularity requirements.

Substrate-specific strand (§§9.3–9.11). The C_6-wheel architecture as a *realisation* of Theorem 9's admissibility, under explicit C_6-boundary-class scope. This is the *construction* result: it exhibits a substrate that produces 3+1.

The two strands together. *Theorem 9 selects 3+1 model-independently; the $K = 7$ wheel architecture provides a minimal candidate mechanism by which a VERSF substrate can realise that selection.* The two strands are logically independent: Theorem 9 does not depend on the $K = 7$ construction, and the $K = 7$ construction does not strengthen the admissibility result beyond realising it. Theorem 9 alone leaves open whether any substrate produces 3+1; the $K = 7$ construction alone could be a coincidence. Together they answer both halves of the dimensional question — *is 3+1 forced?* and *can a substrate produce it?* — closing the loop. The substantive content of the $C_6 \to W_7$ transition is **topological-operational** (extension of the admissible observable algebra via a topologically distinguished reference vertex), not spectral.

Schematic of the $C_6 \to W_7$ transition. A fast intuition anchor before the formal development:

```
     C₆  (vertex-transitive, no reference)            W₇  (hub breaks
transitivity)

         v₁ ─── v₂                                         v₁ ─── v₂
         /       \                                        /│       │\
        v₆        v₃       ──── add hub h ────►        v₆ │   h   │ v₃
         \       /                                        \│       │/
         v₅ ─── v₄                                         v₅ ─── v₄
                                                           (each v_i
connected to h)

         |V| = 6                                           |V| = 7
         dim 𝒪(C₆)  = 5                                    dim 𝒪(W₇)  = 6
         rank(B₆)   = 5                                    rank(W₇)   = 6
         saturated, but no radial reference                saturated, with
radial reference

                        observable-algebra extension: dim(𝒪(W₇)) − dim(𝒪(C₆)) =
1
                        (the new dimension: x_i − h, verified independent by
Lemma 4')
```

The transition is characterised by three simultaneous changes: hub augmentation breaks vertex-transitivity (Lemma 2'), the admissible observable algebra extends by exactly one (radial reference, Lemma 4'), and the constraint operator's rank keeps pace to satisfy saturation (Revised Commitment Principle).

9.2 Theorem 9 (Dimensional Admissibility from Constraint Closure)

Premises: A1, A2, A3; (CGH.0) producing some n-dimensional spatial manifold (n unspecified); BCB conservation; persistence of bound structure; long-range regularity.

Statement. Conditional on (CGH.0) supplying an emergent spatial manifold of some dimension n, any substrate satisfying A1–A4 plus the listed admissibility requirements has $n = 3$. Combined with the temporal-ordering content of A2 (see Step 5 below for the qualifications attaching to this), the emergent structure is 3+1.

Proof.

Step 1 — pre-geometric. A1–A4 reference no dimension. The dimension n on which Steps 2–4 operate is supplied by (CGH.0), which is a standing hypothesis rather than a derivation. The result of Theorem 9 is therefore properly read as: *conditional on coarse-graining producing some spatial manifold, the dimension of that manifold is forced to 3.* Theorem 9 does not establish that coarse-graining produces a manifold; that is (CGH.0).

Step 2 — BCB conservation and its substrate-derived role. By A1, A2, and Definition 3, the record current C^μ satisfies $\partial_\mu C^\mu = s_c$. In the static regime this reduces to $\nabla\cdot\mathbf{J}_{bound} = S$ — a Gauss-law-type conservation statement at the emergent-geometry level. This step carries the substantive VERSF input that distinguishes the present argument from the Tangherlini–Tegmark–Ehrenfest line. Standard derivations take Gauss's law as a postulate of classical electrostatics or gravitation; in VERSF, the Gauss-law form is *derived* from A1, A2 and the record-current structure of Definition 3, which were themselves stipulated pre-geometrically. The Gauss-law scaling then operates on a quantity (bound-record flux) that has a substrate origin rather than a phenomenological one.

Step 3 — Gauss-law scaling. Conserved radial flux through an $(n-1)$-sphere distributes over area $\propto r^{(n-1)}$, giving $\Phi(r) \propto r^{(2-n)}$ for $n \neq 2$ and $\Phi(r) \propto \log r$ for $n = 2$. This step uses the manifold structure supplied by (CGH.0) and the conservation law of Step 2; it does not introduce additional VERSF axioms.

Step 4 — three independent stability arguments converge on n = 3.

(4a) Classical orbital stability. For a central potential $V(r) = -k/r^\alpha$ with angular momentum L, the effective potential $V_{eff}(r) = V(r) + L^2/(2mr^2)$ admits a stable minimum (where $d^2V_{eff}/dr^2 > 0$ at the extremum) only when $\alpha < 2$; direct calculation gives $d^2V_{eff}/dr^2 \propto (2 - \alpha)$ at the extremum. Since gravitational $\alpha = n - 2$ by Step 3, stability requires $\mathbf{n < 4}$.

Within $n \in \{3\}$, Bertrand's theorem distinguishes $n = 3$ sharply: the only central potentials producing closed bound orbits for *every* initial condition compatible with boundedness are the inverse-square ($V \propto -1/r$) and the harmonic oscillator ($V \propto +r^2$) potentials. Of these, the inverse-square is the gravitationally relevant case and corresponds to $\alpha = 1$, hence $n = 3$ via Step 3's relation $\alpha = n - 2$. The harmonic oscillator is irrelevant for an attractive flux-conserving potential (which has $\alpha = n - 2 > 0$ only for $n > 2$). Therefore $n = 3$ is uniquely picked out among admissible $n \in \{3\}$ as supporting closed bound orbits.

(4b) Quantum bound-state stability. The Schrödinger equation in n spatial dimensions with $V = -k/r^{(n-2)}$ has a bounded ground state iff the potential is less singular than $1/r^2$ at the origin (Ehrenfest 1917; Tangherlini 1963; Gurevich–Mostepanenko 1971): $n - 2 < 2$, i.e., $\mathbf{n < 4}$. For $n \geq 4$, the quantum ground state collapses, in violation of A3 (bounded capacity prevents unbounded record concentration at a single point).

(4c) Long-range regularity, tied to A3. For $n = 1$, conserved flux does not dilute (Φ grows linearly); for $n = 2$, $\Phi \propto \log r$ diverges at both limits. The disqualifying property is not the form of the potential per se but its incompatibility with A3. A3 requires bound structures to occupy bounded capacity regions and to be independently realisable — distinct bound configurations cannot be forced to coexist or to fail to localise. Universal confinement in $n = 1$ prevents independent bound structures from existing as separable entities; logarithmic divergence in $n = 2$ prevents either short-range localisation or long-range non-interaction. Both fail A3's independence requirement: $\mathbf{n > 2}$ is required.

Combining: $n \in \{1, 2\}$ excluded by (4c) via A3; $n \geq 4$ excluded by (4a) and (4b) independently. **n = 3 uniquely.**

Step 5 — the temporal dimension from A2. The irreversibility of A2 supplies a partial order on commitment events at the substrate level (the tick-versus-bit distinction: reversible updates versus irreversible records). Under coarse-graining respecting (CGH.0)'s Lorentzian-manifold structure, this partial order is represented by one timelike dimension on $\mathscr{M}$.

We acknowledge that this temporal-dimension argument is substantially thinner than the spatial argument of Step 4. A partial order is not, by itself, a 1-dimensional Lorentzian time direction; the passage from "irreversible commitment ordering" to "global timelike vector field on $\mathscr{M}$" requires further structure beyond A2 alone. The standing hypothesis (CGH.0) supplies the global Lorentzian signature; A2 supplies only the orientation. The full substrate-first construction of the temporal dimension — including arguments for why exactly one temporal dimension rather than zero or two — is the subject of the proto-time and emergent Lorentz invariance companion paper. The present paper takes the +1 as conditional on (CGH.0)'s Lorentzian signature plus A2's orientation, with the substrate-first derivation deferred. *The dimension count 3+1 in this paper is therefore better read as "n = 3 derived; +1 conditional on (CGH.0) + A2."*

Step 6 — covariant completion. With $n = 3$ and one temporal dimension, the covariant-completion argument of Theorem 7 applies and yields GR. ■

What Theorem 9 establishes. This is an *admissibility* result conditional on (CGH.0): given that coarse-graining produces some spatial manifold, the dimension of that manifold is forced to n = 3. The complementary question — does some specific substrate architecture realise such a manifold? — is the substrate-specific strand below.

Added value over the Tangherlini–Tegmark line. Tangherlini (1963) and Tegmark (1997) observed that classical and quantum stability arguments converge on n = 3, taking the relevant scaling laws (Gauss's law, the Schrödinger equation in n dimensions) as ambient givens. The VERSF contribution at Theorem 9 is to embed these arguments within a substrate programme such that the inputs they take as given are themselves substrate-derived: BCB conservation is a consequence of A1, A2, and Definition 3 (Step 2), not a postulate of classical electrostatics; the bound-state stability requirement is tied to A3 rather than imported from atomic physics; long-range regularity at n = 1, 2 is excluded specifically through A3's independence requirement (Step 4c). The classical–quantum stability arguments themselves are not novel; their *recasting as consequences of pre-geometric axioms* is the novel framing. A reader whose interest is the dimensional result alone may find Tangherlini and Tegmark sufficient; a reader interested in why n = 3 should be a *consequence* of the same axioms producing SR and GR will find Theorem 9's embedding within VERSF the substantive contribution.

9.3 The substrate architecture under examination

The VERSF programme's canonical closure cell, derived in the K = 7 No-Go Theorem, is the wheel $W_7 = C_6 + h$: six boundary vertices in a hexagonal cycle, plus one hub vertex connected to every boundary vertex. The closure-constraint count is K = 7 = 6 + 1.

We examine whether this architecture realises 3+1 under coarse-graining. This is **not** an independent derivation of dimensionality. Theorem 9 has already established that 3+1 is the unique admissible dimension; the substrate-specific strand asks only whether the C_6-wheel is consistent with that admissibility — i.e., whether a VERSF substrate can produce the dimension that Theorem 9 forces it to produce.

The scope is the **C_6-boundary class**. Alternative 6-vertex boundary structures exist; §9.10 addresses why C_6 is the VERSF-canonical choice, with explicit attention to what is exhibited (edge-count minimality), what is partially described (fold-interface compatibility), and what is deferred upstream (primitive-eigenvalue normalisation).

9.4 The Revised Commitment Principle and the Operator/Observer Framework

9.4.1 Retraction and replacement of an earlier framing

Earlier formulations of the VERSF commitment mechanism — in v8–v11 of this paper, and in some companion papers — framed commitment as the *elimination of physical null modes* under hub coupling, carrying the language "spectral signature of commitment." This is **inconsistent with what the lemmas actually prove**. Lemmas 2 and 3 establish that B_6 and W_7 each have nullity 1 (the constant vector), purely by graph connectivity. Both null modes are gauge-invariant uniform shifts; both are unobservable to any observer-protocol on their respective cells (Lemma

4). Neither is "physical floating" in a defensible algebraic sense on its own cell. The framing was overloaded and is retracted.

The corrected framing rests on a precise notion of *topological-operational*, which we define here once for use throughout the paper.

Definition (Topological-Operational). A feature of a graph cell G is *topological-operational* if it is (i) invariant under graph isomorphism (graph-topological), and (ii) materially changes the admissible observer-protocol algebra $\mathcal{O}(G)$ (operational). The $C_6 \to W_7$ transition's substantive content is topological-operational in this sense: the hub's introduction is graph-topologically encoded (Lemma 2': automorphism asymmetry, an isomorphism-invariant), and it materially extends $\mathcal{O}(W_7)$ beyond $\mathcal{O}(C_6)$ (operational, via the radial protocol of Definition 6.1's (OP2)).

The corrected framing is then:

Revised Commitment Principle. *Commitment occurs when the constraint operator of a cell saturates the cell's admissible observable algebra. A cell G is committed (relative to its own admissible observables $\mathcal{O}(G)$) when rank(B_G) = dim($\mathcal{O}(G)$).*

Two graph-theoretic quantities — rank of the constraint operator, and dimension of the admissible observable algebra — are directly checkable, and saturation is a precise relation between them. No null-mode language is required.

9.4.2 Operator notation

- **Signed incidence matrix**: edge–vertex matrix B with ±1 entries.
- **Constraint operator**: matrix encoding closure constraints (B_6 for C_6, W_7 for the wheel).
- **Graph Laplacian**: L(G) := B^T B; nullity 1 for any connected G.

9.4.3 Lemma 2' (Automorphism Asymmetry) — foundational

Statement. $\mathrm{Aut}(C_6) = D_6$ acts vertex-transitively; $\mathrm{Aut}(W_7) = D_6$ acts vertex-transitively on the boundary while fixing the hub. Therefore:

(i) **C_6 has no preferred measurement origin.** Without a topologically distinguished vertex, no observer-protocol on C_6 alone selects a vertex as a reference; measurements on C_6 are necessarily relational between adjacent boundary vertices.

(ii) **W_7 has exactly one topologically distinguished vertex (the hub).** The hub is identifiable by vertex degree (hub: 6, boundary: 3) and fixed-point status under $\mathrm{Aut}(W_7)$, enabling observer-protocols that take the hub as a reference origin.

Proof. Direct graph-automorphism computation. ■

Lemma 2' is the foundational graph-theoretic content; everything else in §9.4 follows.

9.4.3a Lemma (No Canonical Reference on Vertex-Transitive Boundaries)

Statement. If G is vertex-transitive — i.e., Aut(G) acts transitively on V(G) — then no observer-protocol on G alone can define a vertex-independent reference observable. Consequently, dim($\mathcal{O}$(G)) excludes radial-reference classes.

Proof. Suppose, for contradiction, that some admissible protocol P on G defines an observable that takes a particular vertex v* as a reference origin. Vertex-transitivity provides an automorphism φ ∈ Aut(G) sending v* to any other vertex w; the same protocol applied via φ would equally well define an observable referenced to w. Since Aut(G) preserves G's edge structure, both observables are admissible by Definition 6.1. The two observables generally disagree on the same state, contradicting the protocol's well-definedness as a single observable. Hence no vertex-independent reference can be selected. ■

Application across the antipodal-paired class. All canonical 6-vertex antipodal-paired graphs are vertex-transitive: C_6 (Aut = D_6 acts transitively), the octahedron $K_{2,2,2}$ (Aut = O_h acts transitively), the prism $C_3 \square K_2$ (Aut acts transitively), $K_{3,3}$ (Aut = $S_3 \wr S_2$ acts transitively on its bipartition classes). By the lemma, none admits a radial-reference observable on its own. Adding a hub vertex (a degree-distinguished vertex coupled to every boundary vertex) breaks vertex-transitivity in any of these cases, enabling radial-reference protocols.

This identifies *hub augmentation* as the operation that enables the new observable class. The argument does not by itself privilege C_6-plus-hub over octahedron-plus-hub or prism-plus-hub; that selection rests on the §9.10 criteria (A) edge-count minimality, (B) fold-interface compatibility, and (C) primitive-eigenvalue normalisation. The lemma's contribution is to upgrade the §9.10 narrative from "C_6 is minimal" to a sharper claim:

Hub augmentation of any vertex-transitive boundary uniquely enables a new class of (radial-reference) observables. Among hub-augmented antipodal-paired 6-vertex boundaries, W_7 is selected by the §9.10 criteria.

9.4.4 Lemmas 2, 3 (connectivity)

Lemma 2. rank(B_6) = 5, nullity(B_6) = 1; ker(B_6) = span{$\mathbf{1}_6$}.

Lemma 3. rank(W_7) = 6, nullity(W_7) = 1; ker(W_7) = span{$\mathbf{1}_7$}.

Connectivity-level facts; supportive but not load-bearing on their own.

9.4.5 Saturation defined

Definition (Saturation). A constraint operator B_G *saturates* the admissible observable algebra $\mathcal{O}$(G) of cell G if rank(B_G) = dim($\mathcal{O}$(G)). Every linearly independent gauge-invariant observable is fixed (up to additive constant) by some closure constraint.

9.4.6 The Observable Algebra: Structure and Admissibility

Before using dim($\mathcal{O}$(G)) as the right-hand side of the saturation criterion, we specify the algebraic structure of $\mathcal{O}$(G).

Definition 6.0 (Observable Algebra Structure). Let G be a graph cell with vertex set V(G) and edge structure E(G). Define

$$\mathcal{O}(G) := \{ O : V(G) \to \mathbb{R} \mid O \text{ arises from an admissible observer-protocol on } G \},$$

with the operations:

- **Linear structure.** For $O_1, O_2 \in \mathcal{O}(G)$ and $a, b \in \mathbb{R}$, the linear combination $aO_1 + bO_2$ is the observable produced by composing protocols and weighting their outputs; this is in $\mathcal{O}$(G) by closure of admissible protocols under linear composition.
- **Equivalence.** Two observables are identified if they produce identical outputs on every state — i.e., quotient by observational indistinguishability.

Under these operations, $\mathcal{O}$(G) is treated as a *finite-dimensional real vector space*. The dimension dim($\mathcal{O}$(G)) is the maximal number of linearly independent admissible observables on G. The observable algebra is finite-dimensional by A1 (finite distinguishability), which ensures that the dimension counts used in the saturation criterion are well-defined.

Restriction to linear observables. $\mathcal{O}$(G) as defined here consists of observables that are linear functionals of vertex amplitude differences — i.e., the protocols of (OP1) and their linear combinations. Nonlinear admissible observables (products, ratios, polynomial combinations) are excluded from $\mathcal{O}$(G) for the purposes of the saturation criterion. This restriction is necessary for dim($\mathcal{O}$(G)) to be a vector-space dimension; without it, dim($\mathcal{O}$(G)) could be infinite (any set of generators closed under polynomials gives a function algebra of unbounded dimension), and the saturation criterion rank(B_G) = dim($\mathcal{O}$(G)) would lose its content. The restriction matches the linear character of constraint operators: B_G acts linearly on state vectors and saturates a linear observable space. Substrate-level commitment in the sense of Definition 7 is therefore a statement about the linear gauge-invariant content of the cell, with nonlinear observables (when they exist as physical quantities) treated as derived rather than primary.

Lemma (Finite Dimensionality from A1). Because A1 enforces finite distinguishability on any finite cell, the set of inequivalent observer-protocol outputs on G — and hence the set of admissible observables modulo observational indistinguishability — is finite. Therefore $\mathcal{O}$(G) is a finite-dimensional real vector space.

Proof. By A1, the cell G admits only finitely many operationally distinguishable physical states. Any admissible observer-protocol must, by (OP1), produce its output from finitely many amplitude differences read off the cell's edge structure; the space of such outputs has cardinality bounded by the number of distinguishable input states, hence finite. The vector space generated by linear combinations of finitely many generators (under the equivalence of Definition 6.0) has finite dimension. ■

Lemma (Gauge Quotient). Let $\mathbb{R}^{|V(G)|}$ be the space of vertex amplitudes on cell G, and let the global U(1) shift act by $\mathbf{x} \mapsto \mathbf{x} + \lambda\mathbf{1}$ for $\lambda \in \mathbb{R}$. Then

$$\mathcal{O}(G) \cong \mathbb{R}^{|V(G)|} / \mathrm{span}\{\mathbf{1}\} \frown \text{(admissible-protocol image)},$$

where the right-hand side denotes the quotient $\mathbb{R}^{|V(G)|}/\mathrm{span}\{\mathbf{1}\}$ restricted to functions generated by admissible observer-protocols on G.

Proof. By (OP1), every admissible protocol reads only amplitude differences and is therefore invariant under the U(1) shift. The map sending admissible observables to their well-defined values on the quotient $\mathbb{R}^{|V(G)|}/\mathrm{span}\{\mathbf{1}\}$ is therefore a linear injection. Surjectivity onto admissible-protocol images is by construction of $\mathcal{O}(G)$. ■

This makes the standard "subtract one for the gauge mode" step mathematically explicit: $\dim(\mathcal{O}(G)) \leq |V(G)| - 1$, with equality iff the admissible-protocol image fills the quotient. For C_6: $\dim(\mathcal{O}(C_6)) = 5 = 6 - 1$ (boundary edge-differences fill the quotient). For W_7: $\dim(\mathcal{O}(W_7)) = 6 = 7 - 1$ (boundary edge-differences plus radial reference fill the quotient; Lemma 4' verifies the radial direction is genuinely new).

For the cells of present interest:

- $\mathcal{O}(C_6)$ has dimension 5, with basis $\{x_i - x_{i+1} : i = 1, \ldots, 5\}$ (boundary edge differences; the sixth is determined by the others as $\Sigma(x_i - x_{i+1}) = 0$ around the cycle).
- $\mathcal{O}(W_7)$ has dimension 6, with basis $\{x_i - x_{i+1} : i = 1, \ldots, 5\} \cup \{x_1 - h\}$ (boundary edge differences plus the radial-reference observable; their independence is established below as Lemma 4').

Definition 6.1 (Admissible Observables via Observer-Protocols). An *observer-protocol* on cell G is a computational procedure executable on G's vertices and edges, using physically realisable measurement primitives, that produces a real-valued output. The admissible observable algebra $\mathcal{O}(G)$ is the vector space defined in Definition 6.0, restricted to outputs realisable on G's vertex set with G's edge structure.

Following the law-closure paper:

(OP1) Observer-protocols read amplitude differences only — between adjacent vertices or composites — never absolute amplitudes.

(OP2) The admissible observable algebra depends on G's graph structure. A cell with a topologically distinguished vertex (in the sense of Lemma 2') admits **radial-reference protocols** — protocols reading boundary-vertex amplitudes against the distinguished vertex — not realisable on a vertex-transitive cell.

Fence-post statement. "Physically realisable measurement primitives" are specified by the law-closure paper *axiomatically, prior to and independently of any notion of commitment*. They are not defined via commitment events, A2-thermodynamic commitment, or Definition 7's structural commitment. This breaks any potential circularity at the deepest level.

Concrete example of a radial protocol on W_7. The protocol $\bar{R}(\mathbf{y}) := (1/6)\, \Sigma\_i\, (x\_i - h)$ measures the average boundary amplitude relative to the hub. It is realisable on W_7 because the hub is the unique distinguished vertex; *not* realisable on C_6 alone, since no vertex on C_6 can be selected as a reference.

9.4.7 Definition 6 (Observable Null Mode), Definition 7 (Commitment), Lemmas 4 and 4'

Definition 6 (Observable Null Mode, relative to ambient cell). Let B be a constraint operator on cell G with state space V, and $\mathcal{O}(G')$ the admissible algebra on cell $G' \supseteq G$. A vector $v \in \ker(B)$ is:

- *physical relative to $\mathcal{O}(G')$* if some $O \in \mathcal{O}(G')$ gives $O(\mathbf{x} + v) \neq O(\mathbf{x})$;
- *gauge relative to $\mathcal{O}(G')$* if $O(\mathbf{x} + v) = O(\mathbf{x})$ for all $O \in \mathcal{O}(G')$.

The classification is *relative*; a null mode of B on G can be gauge with respect to $\mathcal{O}(G)$ and physical with respect to $\mathcal{O}(G')$ for some larger G'.

Definition 7 (Commitment). Cell G is *committed* (relative to its own admissible algebra $\mathcal{O}(G)$) if B_G saturates $\mathcal{O}(G)$, i.e., $\text{rank}(B\_G) = \dim(\mathcal{O}(G))$.

Proposition (Saturation ⇔ Injectivity on Observables). Let $\pi: \mathbb{R}^{|V(G)|} \to \mathcal{O}(G)$ be the projection from vertex amplitudes to admissible observables (well-defined by Lemma (Gauge Quotient)). Then G is committed in the sense of Definition 7 if and only if

$$\text{(I)}\ \ker(B\_G) \subseteq \ker(\pi), \quad \text{and} \quad \text{(II)}\ \dim(\text{im } B\_G) = \dim(\mathcal{O}(G)).$$

Proof. (⟹) Suppose $\text{rank}(B\_G) = \dim(\mathcal{O}(G))$. Condition (II) is immediate from rank = dim(im). For (I): any $v \in \ker(B\_G)$ is a null mode of the constraint operator. Since $\dim(\text{im } B\_G) = \dim(\mathcal{O}(G))$, the constraints already fix every admissible-observable degree of freedom; so v cannot move $\mathcal{O}(G)$, i.e., $v \in \ker(\pi)$. (⟸) Conversely, suppose (I) and (II) hold. (II) gives $\text{rank}(B\_G) = \dim(\mathcal{O}(G))$, which is the saturation condition. ■

This recasts saturation as a purely algebraic statement about how B_G's kernel sits inside π's kernel, removing any narrative ambiguity. A cell is committed exactly when its constraint operator's null modes are gauge-invisible ($\ker(B\_G) \subseteq \ker(\pi)$) and its constraints span the full admissible-observable space ($\dim(\text{im } B\_G) = \dim(\mathcal{O}(G))$).

Lemma 4 (Kernel–Protocol Correspondence). $v \in \ker(B)$ is gauge relative to $\mathcal{O}(G)$ iff every observer-protocol realisable on G alone takes the same value on $\mathbf{x}$ and $\mathbf{x} + v$. *Proof: forward by definition; reverse by Definition 6.1's construction.* ■

Lemma 4' (Radial Observable Independence). On the wheel W_7, the radial observable $\bar{R}: \mathbf{y} \mapsto x\_1 - h$ is linearly independent of the boundary edge-differences $\{E\_i: \mathbf{y} \mapsto x\_i - x\_{i+1}\} : i = 1,$

…, 5}. Consequently, the admissible observable algebra extension $\dim(\mathcal{O}(W_7)) - \dim(\mathcal{O}(C_6)) = 1$ is genuine.

Proof. Suppose, for contradiction, that $\bar{R} = \Sigma_{i=1}^{5} c_i E_i$ for some real coefficients c_i. Evaluating both sides on the state $\mathbf{y} = (0, 0, 0, 0, 0, 0, h_0)$ with $h_0 \neq 0$ (boundary all zero, hub nonzero):

LHS: $\bar{R}(\mathbf{y}) = x_1 - h = 0 - h_0 = -h_0 \neq 0$.

RHS: $\Sigma c_i (x_i - x_{i+1}) = \Sigma c_i (0 - 0) = 0$.

Contradiction. Hence $\bar{R} \notin \mathrm{span}\{E_i\}$; the seventh independent observable is genuinely new. ■

This closes the algebraic gap underlying $\dim(\mathcal{O}(W_7)) = 6$: the count is not asserted by enumeration but verified by linear independence.

9.4.8 Lemma 5 (Antipodal Reduction under Coarse-Graining) — generic

Statement. Let G be a 6-vertex graph admitting an order-2 automorphism $\sigma \in \mathrm{Aut}(G)$ with three orbits of size 2 (three "antipodal pairs"). Under coarse-graining respecting σ and (CGH.2) isotropy, the six directed boundary modes of G reduce to three independent spatial dimensions on $\mathscr{M}$, with independence at the level of observable distinguishability.

Proof. σ acts on the six vertices producing three orbits of size 2; coarse-graining respecting σ and (CGH.2) must identify configurations differing only by σ. The induced equivalence classes are three; under (CGH.2), perturbations along distinct classes produce distinguishable continuum observables. Algebraically: the independent directions form the quotient

$$\text{Independent directions} = \dim(\mathbb{R}^6 / \langle\sigma\text{-identifications}\rangle) = 3,$$

where the quotient is taken in the space of directed boundary modes. Independence is defined in this quotient space under the symmetry σ, not in the raw graph basis (where the six directed modes are linearly dependent in any 2D embedding of C_6). Dimension = 3. ■

Remark on scope. Lemma 5 is **generic** to antipodal-paired 6-vertex graphs:

- C_6 admits the half-turn $r^3 \in D_6$ as σ; pairs are (v_i, v_{i+3}).
- *Octahedron* $K_{2,2,2}$ admits central inversion $i \in O_h$ as σ; pairs are vertex–antipode pairs in 3D.
- *Prism* $C_3 \square K_2$ admits the swap of the two triangular faces as σ; pairs are corresponding vertices in the two triangles.
- $K_{3,3}$ admits an antipodal-pair structure only with a chosen bijection between its bipartition classes; not canonical, marginal.
- K_6 is vertex-transitive without canonical antipodal pairs.

Lemma 5 therefore applies under coarse-graining to C_6, the octahedron, and the prism — all of which produce 3D under (CGH.2) isotropy. **The dimensional realisation 6 → 3 is not C_6-specific.**

This generality has an implication worth stating clearly: the K = 7 architectural strand, if it derives 3+1 at all, derives it via Theorem 9 (admissibility) plus Lemma 5 (which holds for several boundary architectures) plus (CGH.2). It does not derive 3+1 *because* C_6 is the boundary; it derives 3+1 because *any* antipodal-paired 6-vertex boundary plus (CGH.2) suffices. The C_6-specific minimality content lives elsewhere (Lemmas 2', 4, 4'; Theorem 13B), in the observable-algebra extension story, supplemented by the upstream criteria of §9.10.

Consistency with Theorem 9. Lemma 5 provides the coarse-grained realisation of spatial dimensionality consistent with the admissibility constraint of Theorem 9: where Theorem 9 establishes that $n = 3$ is the only admissible spatial dimension under A1–A4 plus stability and regularity, Lemma 5 exhibits the specific coarse-graining mechanism — antipodal reduction under (CGH.2) — by which an antipodal-paired 6-vertex boundary realises that admissible dimension. Theorem 9 forbids $n \neq 3$; Lemma 5 produces $n = 3$ from a substrate-level boundary structure. The two are not redundant: Theorem 9 is dimension-selection at the admissibility level, Lemma 5 is dimension-realisation at the architectural level, and they meet exactly at $n = 3$.

9.5 The A2–Theorem 12 Bridge: A Named Conjecture

A2 is thermodynamic; Definition 7 is structural. Their equivalence is conjectural:

Conjecture (A2–Theorem 12 Bridge). Definition 7's structural commitment is equivalent, up to the Landauer bound, to A2's thermodynamic commitment: the transition from uncommitted to committed configuration requires entropy export $\geq kT \ln 2$ times the log-dimension of the admissible-algebra extension.

The *operational* reading of substrate-architecture results — that $C_6 \to W_7$ thermodynamically implements A2 — remains conjectural. The *structural* and *topological* readings (Theorems 13A, 13B) do not depend on the conjecture.

9.6 Theorem 12 (Topological-Operational Distinction Introduced by Hub Coupling) — auxiliary

Premises: Lemma 2'; Lemma 4; Definitions 6, 6.1, 7; (OC).

Note on role. Theorem 12 is included for completeness — it documents the relational classification of null modes under different observer-protocol algebras — but is *not required for the primary result* (Theorem 13B). A reader uninterested in null-mode classification can skip §9.6 and proceed directly to §9.7 without loss.

Statement. For cells $C_6 \subset W_7$:

(a) *C_6 is committed relative to $\mathcal{O}(C_6)$.* B_6 has rank 5 = dim($\mathcal{O}(C_6)$) (the 5 boundary edge-differences); ker(B_6) is gauge relative to $\mathcal{O}(C_6)$.

(b) *The image of ker(B_6) under the embedding $C_6 \hookrightarrow W_7$ is not contained in ker(W_7), and carries observable content in $\mathcal{O}(W_7)$.* The natural embedding sends $\mathbf{1}_6 \in$ ker(B_6) to ($\mathbf{1}_6$, 0) $\in \mathbb{R}^7$, which is *not* in ker(W_7) (since the hub-coupling constraints x_i − h = 0 require x_i = h, violated by x_i = 1, h = 0). Reading ($\mathbf{1}_6$, 0) against the radial protocol $\bar{R}$ yields a nonzero shift, exhibiting the embedded mode's observable content in $\mathcal{O}(W_7)$.

(c) *W_7 is committed relative to $\mathcal{O}(W_7)$.* W_7's constraint operator has rank 6 = dim($\mathcal{O}(W_7)$) (5 boundary edge-differences + 1 radial class); ker(W_7) is gauge relative to $\mathcal{O}(W_7)$.

Proof. (a) Lemma 4 applied to ker(B_6) under $\mathcal{O}(C_6)$. (b) Direct check: W_7's hub-coupling block contains rows of the form x_i − h, and ($\mathbf{1}_6$, 0) gives 1 − 0 = 1 in those rows, so ($\mathbf{1}_6$, 0) $\notin$ ker(W_7). The radial protocol $\bar{R}$ from §9.4.6 evaluates $\bar{R}((\mathbf{1}_6, 0)) - \bar{R}(\mathbf{0}) = 1$, exhibiting observable content. (c) Direct count, using Lemma 4' for independence: the gauge-invariant degrees of freedom of W_7 under U(1) shift have dimension 7 − 1 = 6 with basis {x_i − x_{i+1}} ∪ {x_1 − h}; W_7's rank-6 constraint operator fixes all six. ■

9.7 Theorem 13B (Primary Result): Observable-Algebra Extension under Hub Coupling

Premises: Lemma 2'; Lemma 4; Lemma 4'; Revised Commitment Principle; the C_6-boundary class.

Theorem 13B. Within the C_6-boundary class, the $C_6 \to W_7$ transition is uniquely characterised by:

(I) **Topological enrichment.** W_7 acquires a topologically distinguished reference vertex (the hub), absent from C_6 — Lemma 2'.

(II) **Observable-algebra extension.** The admissible observable algebra extends from dim($\mathcal{O}(C_6)$) = 5 to dim($\mathcal{O}(W_7)$) = 6, with the new dimension supplied by the radial-reference observer-protocol enabled by (I) and verified independent by Lemma 4'.

(III) **Saturation by the constraint operator.** rank(W_7) = 6 = dim($\mathcal{O}(W_7)$), saturating the extended algebra by the Revised Commitment Principle.

Status: derived under (OC). Primary result of §9.

Restatement. *The core mechanism of commitment is not the rank transition itself but the extension of the admissible observable algebra enabled by the hub. The constraint operator's rank must keep pace with the algebra extension to satisfy saturation; rank-without-extension and extension-without-rank both fail.*

Proof. (I) is Lemma 2'. (II) follows from (OP2) of Definition 6.1, the explicit radial protocol $\bar{R}$, and Lemma 4' (which establishes the radial observable's linear independence from boundary edge-differences). (III) follows by direct count: gauge-invariant degrees of freedom of W_7 are $\{x_i - x_{i+1} : i = 1, ..., 5\} \cup \{x_1 - h\}$, dimension 6 (verified by Lemma 4'); W_7's rank-6 operator fixes all six. ■

Proposition (Degree-of-Freedom Match). For both C_6 and W_7, the gauge-invariant amplitude space dimension equals the admissible-observable algebra dimension:

$$|V(C_6)| - 1 = 5 = \dim(\mathcal{O}(C_6)), \quad |V(W_7)| - 1 = 6 = \dim(\mathcal{O}(W_7)).$$

Proof. The first equality in each line is the Gauge Quotient Lemma applied to the cell. The second equality follows from the explicit basis enumerations: $\mathcal{O}(C_6)$ is spanned by $\{x_i - x_{i+1} : i = 1, \ldots, 5\}$ (the sixth being dependent by cyclic closure); $\mathcal{O}(W_7)$ is additionally spanned by the radial observable (independent by Lemma 4'). ■

This shows that K = 7 is not merely the combinatorial outcome of "six boundary modes plus one hub." It is precisely the vertex count needed for the gauge-invariant amplitude space ($|V| - 1$) to match the admissible observable algebra dimension under hub augmentation: $|V(W_7)| - 1 = 6 = \dim(\mathcal{O}(W_7))$. The constraint operator W_7 then saturates this matched dimension by Lemma 3's rank count, completing the saturation criterion. K = 7 is the smallest connected graph in the C_6-class for which gauge-invariant amplitudes, admissible observables, and constraint rank all coincide at dimension 6.

9.8 Theorems 10 and 11

Theorem 10 (Realisation). Within the C_6-boundary class, W_7 coarse-grains consistently to 3+1. *Proof: three spatial axes from Lemma 5 (generic; under (CGH.2)); one temporal from A2; covariant completion via Theorem 7. The role of W_7 is to realise the dimensional admissibility of Theorem 9, not to derive it independently.* ■

Theorem 11 (Minimality within C_6-class). No connected graph in the C_6-boundary class with $K < 7$ satisfies the admissibility requirements. *Proof: six boundary modes required (Lemma 5); C_6 alone has $\dim(\mathcal{O}(C_6)) = 5 < \dim(\mathcal{O}(W_7)) = 6$ because Lemma 2' shows C_6 has no preferred measurement origin, and Lemma 4' confirms the radial observable is genuinely new; the additional vertex must be independent of the boundary (else breaks antipodal pairing or violates (CGH.2)). $K_{min} = 7$.* ■

9.9 Theorem 13 (Synthesis)

The substantive content of the C_6-wheel architecture is collected as Theorem 13 in three substatements distinguished by what kind of content they carry. They are *co-equal substatements* rather than a sequenced argument: 13B is marked as the primary result not because it is logically downstream of 13A but because it carries the central mechanism of the §9 development. The labelling A/B/C is descriptive (Structural / Topological-Operational / Operational-thermodynamic) rather than ordinal.

Theorem 13A (Structural). W_7 is the minimal C_6-class connected graph supporting three coarse-grained spatial axes under (CGH.2) and the saturation criterion. *Status: derived (within C_6-class).*

Theorem 13B (Topological-Operational — primary). As stated in §9.7. *Status: derived under (OC), within C_6-class.*

Theorem 13C (Operational/thermodynamic). Conjecturally, W_7 is the minimal C_6-class architecture realising A2 thermodynamic commitment at the cell level. *Status: open conjecture, pending §9.5.*

9.10 Why C_6? The Alternative-Boundaries Question

Reading-of-§9 conditionality. All results in this paper labelled "derived within the C_6-boundary class" — Theorems 10, 11, 12, 13A, 13B, 13C — inherit conditionality on the upstream criteria (B) and (C) below being established as exclusions of alternative architectures rather than mere preferences. Criterion (A) is exhibited concretely in this paper; (B) and (C) are upstream-asserted and not re-derived. We assume the upstream conditions hold and flag the dependency mechanically. If (B) or (C) is weaker than asserted upstream, the corresponding §9 results generalise from "C_6-class" to "antipodal-paired 6-vertex class" with edge-count and architectural modifications.

Two readings of the K = 7 wheel. A reader who accepts the upstream selection of C_6 may read the §9 results as *minimal architecture* claims — the wheel W_7 is the smallest architecture realising the topological-operational mechanism within the canonical class. A reader who declines the upstream conditional may read the §9 results as *illustrative architecture* claims — the wheel W_7 is a concrete substrate-level exhibit of the mechanism, with minimality contingent on upstream criteria not re-derived here. Both readings recover the same mathematical content; they differ only in the strength of the existence claim. The paper does not require the stronger reading: even on the illustrative reading, Theorem 13B exhibits a working mechanism by which an antipodal-paired 6-vertex boundary plus a hub realises observable-algebra extension, which is the substantive contribution of §9 independent of uniqueness.

Lemma 5 (§9.4.8) generalises across antipodal-paired 6-vertex graphs. Under coarse-graining with (CGH.2) isotropy, *any* such graph — C_6, octahedron, prism — produces 3D. The choice of C_6 over alternatives is therefore not driven by dimensional realisation. What drives it?

Three criteria are in play in the upstream VERSF programme.

(A) Boundary edge-count minimality. This criterion is exhibited concretely in the present paper. Edge counts of the antipodal-paired 6-vertex graphs:

- **C_6**: 6 edges (the cycle).
- **Prism $C_3 \square K_2$**: 9 edges (two triangles plus three connecting edges).
- **$K_{3,3}$**: 9 edges (complete bipartite on 3+3 vertices; antipodal pairing only with a chosen bijection).

- **Octahedron $K_{2,2,2}$**: 12 edges (every pair of non-antipodal vertices).
- **K_6**: 15 edges.

Among the antipodal-paired class, C_6 uniquely minimises edges. Boundary edge count corresponds in the VERSF substrate to the number of independent boundary closure constraints; minimising it (while preserving the antipodal-pair structure required by (CGH.2)) **favours C_6 as the minimal representative within the antipodal-paired class**. This is a preference under minimality, not a hard exclusion of alternatives; criteria (B) and (C) below are intended to upgrade this from preference to selection at the upstream level.

Bridge to substrate physics. The mapping "graph edges → independent boundary closure constraints" is itself upstream content of the K = 7 No-Go Theorem, not a generic graph-theoretic principle. In the simplicial substrate, each boundary edge of a closure cell carries one closure constraint relating the amplitudes at its endpoints; the boundary edge count therefore equals the constraint count by construction of the substrate's closure structure. This correspondence is asserted upstream, not re-derived here. Without it, criterion (A)'s minimality preference would be a graph-theoretic observation rather than a substrate-physical principle, and the question "why minimise edges?" would lack a substrate-side answer. Readers requiring rigorous establishment of the edge ↔ constraint correspondence should consult the K = 7 No-Go Theorem and fold-density paper.

(B) Fold-interface compatibility. The K = 7 No-Go Theorem and the fold-density paper develop a notion of compatibility between the closure cell's boundary structure and the simplicial fold-interfaces of the VERSF substrate. The criterion, as currently understood, is the following: the boundary's edge-incidence pattern must align with the 1-skeleton of the substrate's hexagonal 2-cells under the simplicial closure relation. Under this alignment condition, the boundary's edges map injectively onto a face-cycle of a substrate 2-cell; closure constraints can then be locally enforced cell-by-cell without additional coupling structure.

C_6 — a 2-regular cycle of length 6 — has edge-incidence matching the 1-skeleton of a hexagonal 2-cell exactly: six edges, six vertices, cyclic incidence. The octahedron's edge-incidence (4-regular, 12 edges, triangulated) matches triangular 2-cells, but the substrate's simplicial structure (per the K = 7 No-Go Theorem) is hexagonal rather than triangular at the relevant scale, so the octahedron fails alignment in this sense. The prism (mixed 3- and 4-cycles) fails for analogous reasons.

We emphasise that this description summarises the alignment condition as we understand it from upstream work; full rigorous establishment is the K = 7 No-Go Theorem's territory and is not re-derived here. A determined referee asking "is (B) actually established as an exclusion, or only as a plausibility argument?" should consult the upstream paper directly. As written, the present paper takes (B) as input; its honesty depends on the upstream argument being rigorous.

(C) Primitive-eigenvalue normalisation. The closure operator of the VERSF substrate has a primitive eigenvalue λ_1 entering as a normalisation constant in the gravitational coupling. The K = 7 No-Go Theorem argues that the closure cell's boundary must produce a specific primitive-eigenvalue value to recover Newton's constant in the coarse-graining. This favours the cycle structure C_6 over the octahedron (whose closure operator has a different spectrum).

Caveat on (C). Normalisation-matching arguments are generically softer than structural exclusions: any boundary structure could in principle be made to produce the right primitive eigenvalue by adjusting other substrate parameters, unless those parameters are independently fixed. We believe, but do not establish here, that the upstream argument fixes the relevant parameters independently and so converts (C) from a parsimony preference into a hard exclusion. If the upstream argument is parsimony-only, (C) is a weaker constraint than (A) and (B), and the C_6-vs-octahedron distinction would rest more heavily on the latter two.

Honest assessment of scope. Criterion (A) is exhibited concretely in this paper as a minimality preference. Criterion (B) is partially described — the alignment condition is named — but its full establishment as an exclusion is upstream. Criterion (C) is asserted upstream and may be parsimony-soft rather than hard-exclusion strong. Theorem 13's "minimality within the C_6-boundary class" is conditional on this upstream selection, with the conditionality flagged in the §8 and §9.12 status tables.

9.11 $K \geq 8$ Uniqueness within the C_6-Class

Within the C_6-class, Theorem 11 establishes $K \geq 7$ but not $K = 7$ uniqueness. Augmentations of W_7 to $K = 8, 9, \ldots$ satisfying admissibility and meaningfully distinct from W_7 are a bounded graph-enumeration problem for the commitment-barrier line.

9.12 Status of the Dimensional Programme

Result	Scope	Conditional on	Status
3+1 admissibility (model-independent)	Any substrate	A1–A4, stability, regularity	**Derived** (Theorem 9)
Antipodal reduction 6 → 3	Antipodal-paired 6-vertex graphs	(CGH.2)	**Derived; generic** (Lemma 5)
Observable algebra structure	General	linearity + observational equivalence	**Supplied** (Definition 6.0)
Radial observable independence	$C_6 \to W_7$	direct linear-algebra	**Derived** (Lemma 4')
$K = 7$ wheel realises 3+1	C_6-class	Theorem 9, Lemma 5, §9.10 (B)+(C)	Derived (Theorem 10)
Observable-algebra extension primary mechanism	$C_6 \to W_7$	(OC), §9.10 (B)+(C)	**Derived** (Theorem 13B)
Structural minimality $K = 7$	C_6-class	Lemmas 5, 11, §9.10 (B)+(C)	Derived (Theorem 13A)

Result	Scope	Conditional on	Status
Topological-operational classification	$C_6 \to W_7$	Lemma 2', (OC), §9.10 (B)+(C)	Derived (Theorem 12, auxiliary)
Edge-count minimality of C_6	Antipodal-paired 6-vertex	direct count	**Exhibited as preference** (this paper)
Fold-interface compatibility favouring C_6	—	upstream K=7 No-Go	Asserted upstream, partially described, not re-derived
Primitive-eigenvalue normalisation	—	upstream K=7 No-Go	Asserted upstream, parsimony-type, may be softer than structural exclusion
Thermodynamic commitment (A2 realisation)	C_6-class	§9.5 conjecture	**Open**
C_6 uniqueness across antipodal-paired class	All such graphs	(A)+(B)+(C)	**Partial: (A) exhibited; (B),(C) deferred**
$K \geq 8$ uniqueness within C_6-class	C_6-class	bounded enumeration	**Open**

---

10. Toward Reducing the Standing Assumptions

The previous sections derive SR, Newtonian gravity, and GR conditional on the standing assumptions (CGH), (EFT), (MARK), (OC) and on the matter-side input WEP. A natural follow-on question is whether some of these standing assumptions can themselves be *reduced* to sharper substrate-level conditions — converting them from broad hypotheses into specific admissibility requirements with clearer falsification paths.

This section sketches three such reductions. Each takes a current standing assumption and replaces it with a precise condition together with an admissibility argument, such that the standing assumption becomes a consequence given the new condition. The reductions do not yet *fully derive* the standing assumptions from A1–A4 alone — each still rests on a residual hypothesis — but they sharpen what remains assumed. We are explicit about this status throughout: these are *strong reductions*, not full derivations, and we label them as Theorems A, B, C, D outside the main Theorem 1–13 sequence to flag them as forward-looking strengthenings rather than core derivations.

The honest framing for §10:

These arguments do not yet remove the standing assumptions, but they reduce them to sharper substrate-level conditions: causal-throughput saturation, record-type quotienting, and single-order commitment comparability.

10.1 Theorem A: Operational Lorentz Invariance from Record-Protocol Equivalence (Strong Conditional Theorem)

Setting. Currently in this paper, (CGH.1) is taken as a standing hypothesis: the substrate propagation speed c (derived as finite by Theorem 1) is asserted to be frame-invariant on $\mathscr{M}$. This invariance, combined with (CGH.2) and (CGH.3), suffices for Theorem 2's Lorentzian kinematics. Theorem A reduces the assertion of frame-invariance to a sharper operational condition.

Theorem A. *Under A1, A4, and the operational condition*

(NSF) No Substrate-Frame Detection. *No admissible observer-protocol on $\mathscr{M}$ — built from committed records via Definition 6.1 — can distinguish uniform motion relative to the substrate update frame,*

the effective transformation group on $\mathscr{M}$ is the Lorentz group, with c (the maximum record-relevant propagation speed of Theorem 1) as the invariant speed.

Remark on (NSF)'s status. (NSF) is not a free admissibility hypothesis; it is an instance of A0 (Observer-Invariant Distinguishability, §2.1) applied to substrate frames, with A1 and A2 supplying the substrate-specific content. The argument is direct. By A0, only observer-invariant distinctions are physical; non-invariant distinctions are gauge. The substrate may carry an update structure that defines a preferred frame at the level of its elementary updates — the void's pre-commitment dynamics. But A1 (finite distinguishability) bounds what any finite cell can resolve, and A2 (irreversible commitment) restricts what counts as a record. Observer-protocols on $\mathscr{M}$ are built from committed records (Definition 6.1, OP1: protocols read amplitude differences between vertices linked by edge structure, with the amplitudes themselves constituting committed-record content). Substrate-update activity *below* the commitment threshold — the regime where the preferred update frame, if it exists, would be operationally accessible — is by A1 and A2 precisely the regime that observer-protocols cannot reach. By A0, distinctions inaccessible to admissible observer-protocols are gauge; therefore the substrate update frame is gauge, and any putative preferred-frame distinction is not physical content.

The chain is **A0 + A1 + A2 + (OC) ⟹ (NSF) ⟹ Lorentz invariance.** A0 supplies the master invariance principle; A1 and A2 establish the operational inaccessibility of the substrate update frame; (OC) supplies the observer-protocol primitives; (NSF) is the resulting gauge classification of frame-relative quantities; and Theorem A's argument then forces Lorentz transformations from causal-cone preservation, homogeneity, isotropy, and group structure.

This sharpens Theorem A's reduction structure: Theorem A reduces (CGH.1) frame-invariance not to (NSF) as a new hypothesis but to A0, A1, A2, and the observer-protocol primitives of (OC) — all of which are already standing in the paper. The residual work for upgrading Theorem A from "strong conditional" to "formal" is to make the A0 + A1 + A2 ⟹ (NSF) implication explicit as a lemma; the proof of that implication is essentially what this Remark sketches.

Proof sketch. A1 + A4 give finite c (Theorem 1). (NSF) asserts that observer-protocols built from committed records cannot distinguish motion relative to the substrate's preferred update frame; by observer-comparability (OC), any such undetectable frame-distinction is gauge rather than physical. Suppose, for contradiction, that some inertial transformation T on $\mathscr{M}$ does *not* preserve the causal cone |x| = ct. Then T transforms the cone to a non-symmetric structure, and the asymmetry would be operationally detectable: an observer could measure the cone in two different frames and infer their relative motion by the cone's distortion, giving a record-detectable distinction between frames. By (NSF), no such record exists, so T must preserve the cone. Together with homogeneity ((CGH.2)), isotropy ((CGH.2)), and group structure ((CGH.3)), causal-cone preservation forces T to be a Lorentz transformation by the standard von Ignatowski argument. ■

Reduction status. Theorem A is a *strong conditional theorem*: (NSF) is no longer a free-floating admissibility condition but a consequence of A1 + A2 acting through the observer-protocol primitives of (OC) (per the Remark above). Theorem A therefore reduces (CGH.1) — frame-invariance of c — not to (NSF) as a new hypothesis but to A1, A2, and the observer-protocol framework of (OC) — all of which are already standing in the paper. (NSF) corresponds to the operational content of (CGH.3) (no preferred substrate frame operationally accessible), expressed in terms checkable against Definition 6.1's primitives rather than asserted as a global geometric statement. The residual hypothesis is empirical: the assertion that admissible protocols genuinely cannot distinguish substrate motion is in principle testable, and corresponds to the precise meaning of "preferred-frame inaccessibility" in analogue-gravity systems (Unruh; Barceló–Liberati–Visser). Full formalisation would inline the A1 + A2 ⟹ (NSF) implication as a lemma; the proof of that implication is essentially what the Remark sketches, and remains the residual work for upgrading Theorem A from "strong conditional" to "formal."

What this does and does not establish. Theorem A does not yet derive Lorentz invariance from A1–A4 alone with full rigour; the A1 + A2 ⟹ (NSF) implication remains at the level of an argument sketch (the Remark) rather than a formalised lemma. What is gained is twofold: first, a precise operational statement of what would have to be true for Lorentz invariance to fail — an admissible observer-protocol that distinguishes substrate motion would constitute a falsification, sharpening the empirical content of the claim; second, the reduction of (CGH.1)'s residual hypothesis to existing axioms (A1, A2) and existing standing assumptions ((OC)), with no new free condition introduced. Theorem A is therefore tighter than its v23 framing suggested: it does not require accepting (NSF) as a new hypothesis but only the formalisation of an implication from existing premises.

10.2 Theorem B: Universal Coupling from Record-Type Quotienting (Strong Reduction)

Setting. Currently, Theorem 5's derivation of Poisson's equation depends on a named-but-not-derived weak equivalence principle: every bound record contributes identically to capacity strain regardless of internal type. Theorem B reduces this assumption to a sharper admissibility condition.

Theorem B. *Suppose the capacity loading functional admits a record-type decomposition*

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

where a labels record type and ρ_a is the density of type-a bound records. Under the operational condition

(NPC) No Primitive Gravitational Charges. *Record types {a} are effective-field labels rather than primitive substrate labels: no admissible observer-protocol independently detects a as a substrate-level charge,*

together with operational non-redundancy from (OC), the coefficients κ_a must be equal: $\kappa_a = \kappa$ for all a. Hence $\Sigma(x) = \kappa \Sigma_a \rho_a(x) = \kappa \rho_{bound}(x)$, and gravitational coupling is universal across bound record types.

Proof sketch. Suppose, for contradiction, that $\kappa_a \neq \kappa_b$ for some pair of record types a, b. Then gravity itself — through capacity strain — distinguishes type-a configurations from type-b configurations, even when their bound-record counts are equal. This produces a substrate-level distinction between a and b that is gravitationally detectable. By (NPC), record types are effective-field labels rather than primitive substrate labels, so a substrate-level distinction between a and b would constitute additional substrate structure not present in the commitment-density description. By operational non-redundancy from (OC), additional substrate structure must be independently observable — but (NPC) explicitly excludes independent observability of record-type labels at the substrate level. The structure must therefore be quotiented out, forcing $\kappa_a = \kappa_b$. Iterating across all pairs: $\kappa_a = \kappa$ for all a. ■

Remark on (NPC)'s status. Like (NSF) and (GC), (NPC) is not introduced freely. It is an instance of A0 (Observer-Invariant Distinguishability, §2.1) applied to record-type distinctions, with the fold-density paper supplying the substrate-specific content that the κ_a coefficients refer to. The argument is structural: by A0, distinctions between record types are physical only if admissible observer-protocols can independently resolve them. If all bound matter is composed of fold configurations (fold-density paper), and the fold carries no observable labels beyond commitment density, then by construction admissible protocols cannot independently resolve type-a from type-b at the substrate level — the type label is an effective-field artefact of the configuration, not a primitive substrate quantum number. By A0, distinctions that admissible protocols cannot resolve are gauge; therefore differences between κ_a coefficients would be A0-gauge differences, not physical ones. Forcing equality:

If all matter is composed of configurations of a single fold, and the fold carries no additional observable labels beyond commitment density, then gravity couples universally to bound matter.

The two conjuncts of this statement do separate work. The first — *all matter is fold configurations* — is the reductionist content of the fold-density paper: distinct effective-field record types (electric, weak, colour, mass-bearing, etc.) are different configurations of the same primitive substrate degree of freedom, not different primitive substrate degrees of freedom. The second — *the fold carries no observable labels beyond commitment density* — asserts that admissible observer-protocols on fold configurations distinguish only commitment density, not any additional internal-fold quantum number. Together, these two claims combine with A0 to entail (NPC): the {a} of the record-type decomposition above are effective labels of fold configurations, not primitive substrate charges, and gravitational distinctions between them are gauge by A0.

The chain is **A0 + fold-density-paper substrate content ⟹ (NPC) ⟹ universal coupling.** A0 supplies the master invariance principle; the fold-density paper supplies the single-fold reductionist content; (NPC) is the resulting gauge classification of record-type distinctions; and Theorem B's argument then forces $\kappa_a = \kappa$.

Status of (NPC), honestly assessed. (NPC) inherits its status from this fold-substrate content combined with A0. Three observations:

1. (NPC) is therefore *grounded*, not free: it follows from a definite substrate claim (fold-density paper) plus the master invariance principle (A0) rather than being introduced as a new admissibility hypothesis. The structure of Theorem B's reduction matches Theorems A and C in this respect.
2. The grounding combines an A0-internal step (A0 ⟹ record-type distinctions are gauge if not independently observable) with an upstream-content step (fold-density paper ⟹ record types are not independently observable). The first step is internal to A0; the second is upstream programme work. Parallel to §9.10's criteria (B) and (C), the present paper takes the upstream content as input. The honest framing: Theorem B reduces WEP to (NPC), and (NPC) reduces in turn to A0 + fold-density-paper substrate content. The full chain is WEP ⟸ (NPC) ⟸ A0 + fold-substrate content.
3. The falsification path is unchanged: WEP violations occur if and only if the substrate carries primitive record-type charges beyond commitment density. The empirical signature would be composition-dependent free-fall, the content of fifth-force / Eötvös-style tests. What changes in v30 is how (NPC) is justified — as an A0-instance combined with fold-substrate content rather than as a free hypothesis or as fold-substrate alone — but the empirical content of the prediction is the same.

Reduction status. Theorem B is therefore a *strong reduction* whose residual hypothesis is no longer (NPC) read as an independent admissibility condition but the upstream fold-density-paper claim about single-fold substrate combined with A0. Full formalisation has two pieces: (i) inline the fold-density paper's single-fold substrate claim with appropriate scoping, parallel to §9.10's treatment of upstream criteria; (ii) prove A0 + fold-content ⟹ (NPC) as a lemma, parallel to the residual lemmata for Theorems A and C.

Predicted falsification path (unchanged). WEP violations occur if and only if the substrate contains additional primitive record-type charges — quantum numbers intrinsic to the substrate's commitment structure rather than emerging as effective-field labels of fold configurations. The Standard Model's gauge charges (electric, weak, colour) are effective-field labels in the relevant sense, and Theorem B predicts WEP holds across them. A primitive "gravitational charge" distinct from mass-energy — not arising from gauge-theoretic effective-field structure on fold configurations — would violate WEP and would be observable as composition-dependent free-fall. Eötvös-style experimental tests place tight bounds on such primitive charges.

What this does and does not establish. Theorem B does not derive WEP from A1–A4 alone; the residual hypothesis is now the fold-density-paper substrate content rather than (NPC) as a free condition. What is gained is a precise structural reading of WEP-violation: the substrate would have to carry primitive charges beyond fold commitment-density. The reduction converts WEP from a phenomenological postulate into a consequence of fold-substrate non-redundancy, with the residual hypothesis localised to the upstream programme claim about substrate primitives.

10.3 Theorem C: Single Time Dimension from Commitment-Order Consistency (Plausible Theorem Requiring Formalisation)

Setting. Theorem 9 Step 5 currently supplies the temporal dimension via "A2's commitment ordering supplies one timelike dimension on $\mathscr{M}$", with the v21 acknowledgment that this is substantially thinner than the spatial argument. Theorem C addresses this thinning by giving a real admissibility argument for *exactly one* temporal dimension.

Theorem C. *Suppose commitment events on the substrate form an irreversible partial order ($\prec$) by A2, and that this order satisfies*

(GC) Global Comparability. *If $e_i \prec e_j$ and $e_j \prec e_k$, then $e_i \prec e_k$ (transitivity), and admissible observers agree on the order of any two comparable events.*

Then the minimal admissible coarse-grained representation of $\prec$ has exactly one timelike dimension.

Remark on (GC)'s status. (GC) is not an independent admissibility condition; it is an instance of A0 (Observer-Invariant Distinguishability, §2.1) applied to event orderings, with A2 supplying the irreversibility content and (OC) supplying the observer-protocol primitives. By A0, only observer-invariant distinctions are physical; non-invariant distinctions are gauge. If two observers reading the same pair of comparable commitment events disagreed on their order, the temporal-order distinction would be observer-relative — by A0, gauge — and so the events would not be facts in the operational sense the programme requires. The Facthood = Invariant Comparability Proposition (§2.1) is the commitment-event-specific application of A0 that crystallises this: facts are commitments invariant under all admissible observer-protocol comparisons.

The requirement of global comparability is therefore not an independent hypothesis but a consequence of what A0 forces commitment events to satisfy in order to be facts. A commitment whose ordering depends on the observer is, by A0, an observer-dependent commitment — not a fact. Therefore, the existence of facts entails global comparability, and this in turn enforces a single consistent temporal ordering at the coarse-grained level. The argument chain is:

A0 + A2 + (OC) $\Longrightarrow$ Facts exist $\Longrightarrow$ (GC) holds $\Longrightarrow$ exactly one temporal dimension.

This connects five programme components: the master invariance principle (A0), thermodynamics (A2's irreversibility supplying the order), observer theory ((OC) supplying observer-protocol comparability), ontology (what counts as a fact, §2.1), and geometry (the temporal dimension itself). The clean reading: *time exists because facts exist, and facts exist because A0 forces only invariant distinctions to be physical.* Without A0, observer-relative distinctions might count as facts and the chain would not close at one dimension; with A0, observer-relative distinctions are gauge by definition, facts are exactly the invariant commitments, and the minimal coarse-grained representation has exactly one temporal dimension.

(GC) reduces accordingly. At the proof level, (GC) follows from (OC) + A2: A2 supplies the order, (OC) supplies its observer-invariant character. At the foundational level, (GC) is a definitional consequence of A0 + facthood — observer-relative commitments are gauge by A0, so the substrate's facts (the non-gauge commitments) are globally comparable by construction. We list (GC) explicitly because it is the operational content actually used in the proof, but it inherits its status from (OC) + A2 at the derivational level and from A0 + the Facthood = Invariant Comparability proposition at the constitutive level.

Proof sketch. Three cases.

Zero temporal dimensions. A coarse-grained structure with no temporal dimension cannot represent the asymmetric relation $\prec$. Any pair (e_i, e_j) related by $\prec$ would have to be assigned a single coordinate value or a symmetric distance, neither of which encodes the irreversible ordering required by A2. *Excluded.*

One temporal dimension. A single temporal coordinate τ admits the representation $e_i \prec e_j \Leftrightarrow \tau(e_i) < \tau(e_j)$. The order is faithfully represented; (GC)'s transitivity and admissibility are automatic. *Admissible.*

Two or more independent temporal dimensions. Suppose we have two timelike coordinates τ^1, τ^2. For two events e_i, e_j with $\tau^1(e_i) < \tau^1(e_j)$ but $\tau^2(e_i) > \tau^2(e_j)$, the temporal-order relation is direction-dependent: along τ^1, $e_i \prec e_j$; along τ^2, $e_j \prec e_i$. By (GC), admissible observers must agree on event order, so this configuration is inadmissible. The only configurations consistent with (GC) are those where every pair is ordered identically along all temporal directions, which collapses the multi-dimensional temporal structure to its one-dimensional projection. *Reduces to the single-dimensional case.*

Therefore, the minimal admissible coarse-grained temporal structure has exactly one dimension. ■

Important caveat on metric structure. Theorem C derives the *count* (one timelike dimension), not the *signature* (Lorentzian rather than Galilean or Carrollian). The Lorentzian signature combining the temporal dimension with the spatial 3-manifold is supplied by (CGH.0)'s standing hypothesis of an emergent Lorentzian manifold and by Theorem 2's derivation of Lorentz transformations under (CGH). Theorem C contributes the dimension count; (CGH.0) and Theorem 2 contribute the metric structure.

Reduction status. Theorem C is a *plausible theorem requiring formalisation*: the proof sketch is sound and (GC) is no longer a free-floating admissibility condition but a consequence of (OC) plus A2 (per the Remark above). Theorem C therefore reduces "the +1 from A2" not to (GC) as a new hypothesis but to (OC) and A2 themselves — both of which are already standing in the paper. At the foundational level, (GC) is further reduced via the chain *facts exist ⟹ (GC) ⟹ one time dimension* (§2.1 Proposition: Facthood = Invariant Comparability): observer-relative commitments are not facts, so the substrate's facts have globally comparable orderings by definition. The two readings work in concert — proof-level via (OC) + A2 ⟹ (GC), foundational-level via the Facthood = Invariant Comparability chain — and either suffices for Theorem C's argument.

The corresponding falsification path is sharper than v25's framing. A "non-globally-comparable commitment ordering" would not be a deficient fact: by the Facthood = Invariant Comparability proposition, it would belong to the category of *observer-dependent commitments*, distinct from the category of *facts*. Theorem C therefore predicts: if any future observation produces commitment events that admissible observers genuinely disagree about in temporal order, those events are observer-dependent commitments rather than facts; the substrate's actual facts continue to admit one timelike dimension by the chain above. The empirical test reduces to whether the disputed events can be reclassified as observer-dependent (in which case Theorem C is preserved) or whether they must be facts (in which case (OC) is in question and the entire programme's foundation is in doubt). The hard boundary set by §2.1's proposition makes this binary structure operational.

Full formalisation would inline both implications — (OC) + A2 ⟹ (GC) and Facthood = Invariant Comparability ⟹ (GC) — as numbered lemmata; the proofs of those implications are essentially what the Remark above sketches, and remain the residual work for upgrading Theorem C from "plausible" to "formal."

10.4 Theorem D: Dimensional Closure (Combination)

Theorem D. *Combining Theorem 9 (spatial admissibility) with Theorem C (temporal admissibility), the only admissible coarse-grained dimensional count under A1–A4, BCB conservation, finite-capacity stability, observer-comparability, and commitment-order consistency is*

$$D = 3 + 1.$$

Status: *follows from Theorems 9 and C in their respective conditional regimes.*

Important scope statement. Theorem D fixes the *dimension count* 3+1; it does not by itself supply the smooth-manifold or Lorentzian-metric structure on which §§4–7 work. That structure remains a (CGH.0) standing hypothesis, with Theorem 2's Lorentzian kinematics emerging from (CGH.1)+(CGH.2)+(CGH.3). The relation between Theorem D and (CGH.0) is therefore: Theorem D removes the dimension-count ambiguity from (CGH.0), reducing what (CGH.0) must hypothesise to "an emergent smooth Lorentzian manifold of *some* dimension" — with the dimension then fixed at 3+1 by Theorem D.

10.5 Status of the strengthenings

Theorem	Reduces	To	Status
Theorem A	(CGH.1) frame-invariance of c	(NSF) — itself a consequence of **A0** + A1 + A2 + (OC)	**Strong conditional theorem** (residual work: prove A0 + A1 + A2 ⟹ (NSF))
Theorem B	WEP / universal scalar sourcing	(NPC) — itself a consequence of **A0** + fold-density-paper single-fold substrate content	**Strong reduction** (residual work: inline fold-substrate content + prove A0 + fold-content ⟹ (NPC))
Theorem C	"+1 from A2" (Theorem 9 Step 5)	(GC) — itself a consequence of **A0** + (OC) + A2 via the chain *A0 + A2 + (OC) ⟹ Facts exist ⟹ (GC) ⟹ 1 time dimension*	**Plausible theorem requiring formalisation** (residual work: prove (OC) + A2 ⟹ (GC); inline the A0-Facthood chain)
Theorem D	(CGH.0) dimension ambiguity	Theorems 9 + C	**Derived** combination, scope-limited to dimension count

Unifying structure. A0 (Observer-Invariant Distinguishability) is the master principle generating (NSF), (NPC), (GC) when combined with substrate-specific content (A1+A2 for NSF; fold-density-paper for NPC; A2+OC+Facthood for GC). Each of NSF, NPC, GC is a separate *application* of A0 to a specific candidate distinction — substrate frames, record-type labels, event orderings — but the operative principle is one and the same. The §10 architecture is therefore: A0 supplies the criterion of physicality; A1–A4 + (OC) + upstream content supply the substrate-specific facts to which A0 is applied; the resulting gauge classifications produce NSF, NPC, GC; and these in turn produce Lorentz invariance, universal coupling, and the temporal dimension. *§10's three reductions are not three independent strengthenings but three instances of one master principle.*

10.6 What §10 establishes and what it does not

Establishes:

- That (CGH.1), WEP, and the temporal-dimension argument can each be reduced to sharper substrate-level admissibility conditions.
- A precise operational meaning for each reduction: (NSF) for Lorentz, (NPC) for WEP, (GC) for temporal dimension.
- A combined dimensional-closure result (Theorem D) fixing D = 3+1 within the reduction structure.
- Concrete falsification paths: substrate-frame-detecting protocols would falsify Theorem A; primitive substrate gravitational charges (i.e., labels on fold configurations beyond commitment density) would falsify Theorem B; non-globally-comparable commitment orderings would falsify Theorem C.
- For all three theorems, *grounding* of the residual condition rather than free introduction. (NSF) follows from A0 + A1 + A2 + (OC); (NPC) follows from A0 + fold-density-paper

substrate content; (GC) follows from A0 + (OC) + A2 via the Facthood = Invariant Comparability proposition. None of (NSF), (NPC), (GC) is introduced as a new admissibility hypothesis; each is sourced to A0 plus definite substrate-specific content.

- The unifying structural insight: **(NSF), (NPC), and (GC) are not three independent admissibility conditions but three applications of A0 (Observer-Invariant Distinguishability) to substrate frames, record-type labels, and event orderings respectively.** A0 is the master principle; A1–A4, (OC), and upstream programme content supply the specific facts to which A0 is applied. The §10 architecture is therefore top-down rather than three-parallel: A0 + A1–A4 ⟹ (NSF) + (NPC) + (GC) ⟹ Lorentz + WEP + 1-time-dimension ⟹ SR + GR + 3+1.

Does not establish:

- Full formal derivation of any of (CGH.1), WEP, or the temporal dimension from A1–A4 alone. The three reductions inherit residual work of two distinct kinds:
  - **Internal formalisation** for Theorems A and C: the implications A1 + A2 ⟹ (NSF) and (OC) + A2 ⟹ (GC) need to be promoted from prose-level Remarks to numbered lemmata.
  - **Upstream-content formalisation** for Theorem B: the implication (fold-density-paper substrate content) ⟹ (NPC) needs to be made explicit, and the upstream fold-density-paper claim itself needs scoping parallel to §9.10's treatment of upstream criteria.
- (CGH.0)'s smooth-manifold structure. Theorem D fixes the dimension count but not the metric or smoothness.
- The (EFT) and (MARK) standing hypotheses, which are not addressed in §10.

The §10 reductions are best understood as **strong steps toward removing the remaining standing assumptions, with each reduction localising what residual condition would still need substrate-first derivation, and with A0 as the unifying master principle that generates all three.** Theorems A and C ground to A0 plus existing axioms (A1, A2) and standing assumptions ((OC)) without requiring upstream programme content; their residual work is internal formalisation. Theorem B grounds to A0 plus the upstream fold-density paper's substrate content, parallel to §9.10's grounding of C_6 in the upstream K = 7 No-Go Theorem. The structural asymmetry is therefore not "A/C are grounded vs B is free" — all three are grounded — but "A/C ground to existing internal premises while B grounds to upstream programme content," with A0 supplying the master principle in all three cases. The full programme target — deriving SR + GR from A0 + A1–A4 with no residual conditions — remains open, but §10's reductions chain the remaining work to either internal formalisation (A, C) or to upstream programme work already in progress (B). Each chain is now visible, and each is an application of A0.

---

11. Conclusion

This paper has derived Special Relativity, Newtonian gravity, and General Relativity as the effective physics of record dynamics on a 3+1-dimensional Lorentzian manifold, conditional on four standing assumptions: (CGH), (EFT), (MARK), and (OC) — the last implemented through the imported observer-protocol framework with axiomatic primitives.

The dimensional programme has two complementary halves. Theorem 9 establishes model-independently that 3+1 is the unique admissible spatial dimension under A1–A4 plus standard stability and regularity requirements. The K = 7 wheel architecture provides a minimal candidate mechanism by which a VERSF substrate can realise that selection — exhibiting a candidate Layer-1 cell whose coarse-graining produces 3+1 consistently with what Theorem 9 forces. The two strands are logically independent: Theorem 9 does not depend on the K = 7 construction, and the K = 7 construction does not strengthen the admissibility result beyond realising it. Neither pillar is auxiliary: Theorem 9 alone leaves open whether any substrate can produce 3+1, and the K = 7 construction alone could be a coincidence; together they answer both the existence and the realisation question. *Theorem 9 selects 3+1; K = 7 provides a minimal candidate mechanism for realising it.*

Three corrections to earlier framings:

(i) **Commitment is saturation of the admissible observable algebra.** The earlier "elimination of null modes" framing is retracted. The Revised Commitment Principle (§9.4.1) supersedes it: rank(B_G) = dim($\mathcal{O}$(G)).

(ii) **The $C_6 \to W_7$ transition is topological-operational.** The hub introduces a topologically distinguished reference vertex (Lemma 2'), enabling new observer-protocols (Definition 6.1) that extend the admissible observable algebra (Lemma 4'). The rank/nullity content of Lemmas 2 and 3 is supportive but not the primary mechanism. Earlier "spectral signature of commitment" or "physical-to-gauge reclassification" framings are retracted.

(iii) **The K = 7 architecture realises 3+1 admissibility but does not derive 3+1 in isolation.** Dimensional admissibility is selected by Theorem 9 plus (CGH.2). The K = 7 construction is the architectural mechanism by which this admissibility is realised, conditional on the upstream selection of C_6 as the canonical boundary among antipodal-paired 6-vertex graphs (octahedron, prism, $K_{3,3}$-with-pairing). Lemma 5 is generic: its 6 → 3 reduction holds for any antipodal-paired 6-vertex graph. The C_6-specific content is not dimensional reduction but **boundary edge-count minimality** (C_6: 6 edges; alternatives: 9 or more), supplemented by upstream criteria (fold-interface compatibility, primitive-eigenvalue normalisation) that this paper does not re-derive.

Nature of the contribution. VERSF does not re-derive Lovelock or Deser–Feynman. It supplies the pre-geometric substrate inputs those standard arguments require — Lorentz-invariant propagation, universal scalar sourcing, D = 4 emergent dimensionality, universal stress-energy coupling — from A1–A4.

Open items honestly flagged.

- The A2–Definition 7 thermodynamic bridge (§9.5) is a named conjecture.

- C_6 uniqueness across the antipodal-paired 6-vertex class is partially exhibited (edge-count minimality, criterion A) and partially deferred upstream (fold-interface, primitive eigenvalue — criteria B, C) — §9.10. All §9 results labelled "derived within the C_6-class" inherit conditionality on the upstream criteria (B) and (C) being established as exclusions rather than mere preferences.
- $K \geq 8$ uniqueness within the C_6-class is open — §9.11.
- (CGH) has analogue-gravity precedent but no substrate-first proof. §10's Theorem A reduces (CGH.1) to (NSF) No Substrate-Frame Detection, which itself follows from A0 + A1 + A2 + (OC); full formalisation of Theorem A reduces to formalising the A0 + A1 + A2 ⟹ (NSF) implication.
- WEP is named, not derived, in Theorem 5. §10's Theorem B reduces WEP to (NPC) No Primitive Gravitational Charges, which itself follows from A0 + the upstream fold-density paper's single-fold substrate content; full formalisation of Theorem B reduces to inlining the upstream substrate claim with appropriate scoping (parallel to §9.10's treatment of upstream criteria) and formalising the A0 + fold-content ⟹ (NPC) implication.
- The temporal-dimension argument of Theorem 9 Step 5 is sharpened in §10's Theorem C via the chain *A0 + A2 + (OC) ⟹ Facts exist ⟹ (GC) ⟹ one time dimension*. (GC) follows from (OC) + A2 at the proof level, and from A0 + the §2.1 Proposition (Facthood = Invariant Comparability) at the foundational level. The clean reading is *time exists because facts exist, and facts exist because A0 forces only invariant distinctions to be physical*. Full formalisation of Theorem C reduces to formalising the (OC) + A2 ⟹ (GC) implication and inlining the A0-Facthood chain.
- (OC)'s observer-protocol primitives are imported axiomatically from the law-closure paper; substrate-first derivation from A1–A4 is open.
- A full audit of the $K = 7$ No-Go Theorem against the Revised Commitment Principle has not been completed.
- Numerical values of G, Λ, μ are deferred to companion papers.

On the strengthenings of §10. Theorems A, B, C, D each reduce a residual standing assumption to a sharper substrate-level condition with concrete falsification paths, but none yet removes the assumption entirely. (NSF), (NPC), (GC) remain hypotheses, although each is sharper and more specific than the standing assumption it replaces. The full programme target — deriving SR + GR from A1–A4 alone, with no residual standing assumptions — remains open; §10 narrows the gap.

On the master principle A0. The §10 reductions are unified by a single principle made explicit in §2.1: A0 (Observer-Invariant Distinguishability). A0 is not an arbitrary axiom but the minimal consistency requirement for a notion of physical fact to be well-defined: if non-invariant distinctions were allowed, physics would depend on observer-relative descriptions rather than invariant content. A0 states that physical content depends only on distinctions invariant under all admissible observer-protocols; non-invariant distinctions are gauge and do not constitute physical facts. Each of (NSF), (NPC), (GC) is an application of A0 to a specific candidate distinction:

- *(NSF)*: substrate update frames are A0-gauge because A1 and A2 prevent observer-protocols from accessing them.
- *(NPC)*: record-type labels are A0-gauge because the fold-density paper establishes them as effective-field artefacts of fold configurations rather than primitive substrate distinctions.
- *(GC)*: temporal-order disagreements between admissible observers are A0-gauge by definition; the substrate's facts are exactly the commitment events for which order is observer-invariant (§2.1 Facthood Proposition).

The unifying reading: *only observer-invariant distinctions are physical; everything else — space, time, gravity — is what that principle, combined with the substrate dynamics of A1–A4, forces the world to look like.* The architecture is therefore top-down: A0 + A1–A4 + standing assumptions (with each standing assumption itself reducible to A0 + substrate content via §10) ⟹ relativity, gravitation, and dimensionality.

The architectural insight in one line. What the framework ultimately says is this:

Physics = invariant structure of the observable algebra.

Spacetime, gravity, motion, and matter are what remains after quotienting out non-invariant distinctions. The substrate produces a finite-dimensional observable algebra (Definition 6.0; Lemma Finite Dimensionality from A1) whose admissible observer-protocols define an equivalence relation on substrate states; A0 declares the equivalence classes to be the physical content; A1–A4 supply the substrate dynamics whose coarse-graining produces the observed structure. The Lorentz group emerges as the symmetry of substrate frames quotiented out by A0; universal gravitational coupling emerges as record-type labels quotiented out by A0 (under fold-substrate content); the temporal dimension emerges as the minimal coarse-grained representation of irreversible commitment ordering with observer-relative orderings quotiented out by A0. In this sense the entire framework is the elaboration of a single insight: physics is what survives the quotient.

Layered architecture:

- **Layer 1 (Substrate).** A0 (master principle) + A1–A4 (substrate axioms); pre-geometric.
- **Layer 2 (Emergent geometry).** 3+1 admissibility model-independent (Theorem 9); substrate-architectural realisation via the C_6-wheel under explicit scope caveats and the topological-operational mechanism of (ii).
- **Layer 3 (Effective physics).** SR, Newtonian gravity, GR.

Conditional on A0 (Observer-Invariant Distinguishability), the standing hypotheses, and the upstream selection of C_6 as the VERSF-canonical boundary, Special Relativity, Newtonian gravity, and General Relativity are derived as the effective physics of record dynamics on a 3+1-dimensional Lorentzian manifold whose dimensionality is selected model-independently from substrate-level stability and regularity requirements. The C_6-wheel architecture provides a minimal candidate mechanism for realising this admissibility through the topological-operational mechanism of observable-algebra extension — Theorem 9 selects 3+1, K = 7 *provides a minimal candidate mechanism for realising it — with named open questions about thermodynamic interpretation, alternative architectures, and the upstream criteria privileging C_6 over its non-edge-minimal antipodal-paired alternatives. A0 supplies the unifying principle behind the §10 reductions of (CGH.1), WEP, and the temporal-dimension argument: only observer-invariant distinctions are physical; relativity, gravitation, and the 3+1 dimensional structure are what that principle, applied to the substrate of A1–A4, forces the coarse-grained world to look like.*

References

Foundational and historical works on dimensional and stability arguments

Bertrand, J. (1873). *Théorème relatif au mouvement d'un point attiré vers un centre fixe*. Comptes Rendus de l'Académie des Sciences, **77**, 849–853. [Bertrand's theorem on closed orbits, used in §9.2 Step 4(a) to single out n = 3 as supporting closed bound orbits for an inverse-square central force.]

Ehrenfest, P. (1917). *In what way does it become manifest in the fundamental laws of physics that space has three dimensions?* Proceedings of the Amsterdam Academy, **20**, 200–209. [Earliest systematic argument that quantum bound-state stability requires n = 3; cited in §9.2 Step 4(b).]

Tangherlini, F. R. (1963). *Schwarzschild field in n dimensions and the dimensionality of space problem*. Il Nuovo Cimento, **27**, 636–651. [Generalisation of orbital and bound-state stability arguments to n spatial dimensions; cited in §9.2 Step 4(b) and in the §9.2 added-value remark.]

Gurevich, L. E., & Mostepanenko, V. M. (1971). *On the existence of atoms in n-dimensional space*. Physics Letters A, **35**(3), 201–202. [Quantum stability of bound states in n dimensions; cited in §9.2 Step 4(b).]

Tegmark, M. (1997). *On the dimensionality of spacetime*. Classical and Quantum Gravity, **14**(4), L69–L75. [Modern overview of the n = 3 stability arguments and their consilience; cited in §1.3 and the §9.2 added-value remark.]

Analogue gravity and emergent Lorentz invariance

Unruh, W. G. (1981). *Experimental black-hole evaporation?* Physical Review Letters, **46**(21), 1351–1353. [Foundational paper on emergent Lorentzian metric structure in Bose–Einstein condensate analogues; cited in §2.3 as physical precedent for (CGH) and in §10.1's Reduction status remark.]

Barceló, C., Liberati, S., & Visser, M. (2011). *Analogue gravity*. Living Reviews in Relativity, **14**, Article 3. [Comprehensive review of analogue-gravity systems showing Lorentzian causal structure emerging on top of a non-Lorentzian substrate; cited in §2.3 and §10.1 as the established physical precedent for the substrate-emergent-geometry framing.]

Foundations of general relativity

Lovelock, D. (1971). *The Einstein tensor and its generalizations*. Journal of Mathematical Physics, **12**(3), 498–501. [The Lovelock theorem characterising the Einstein tensor as the unique second-order divergence-free tensor in D = 4; used in Theorem 7 Step 1.]

Lovelock, D. (1972). *The four-dimensionality of space and the Einstein tensor*. Journal of Mathematical Physics, **13**(6), 874–876. [Companion result on dimensional dependence of the uniqueness theorem; relevant to Theorem 7's D = 4 assumption.]

Deser, S. (1970). *Self-interaction and gauge invariance*. General Relativity and Gravitation, **1**(1), 9–18. [The bootstrap argument deriving general relativity from massless spin-2 self-coupling consistency; used in Theorem 7 Step 2.]

Feynman, R. P., Morinigo, F. B., & Wagner, W. G. (1995). *Feynman lectures on gravitation*. Addison-Wesley. [Lecture notes presenting the spin-2 bootstrap argument independently of Deser; the "Deser–Feynman" framing references both treatments.]

Tests of relativity and the equivalence principle

Bertotti, B., Iess, L., & Tortora, P. (2003). *A test of general relativity using radio links with the Cassini spacecraft*. Nature, **425**, 374–376. [Cassini Shapiro-delay test constraining PPN-γ; referenced in §5's PPN remark.]

Williams, J. G., Turyshev, S. G., & Boggs, D. H. (2012). *Lunar laser ranging tests of the equivalence principle*. Classical and Quantum Gravity, **29**(18), 184004. [Lunar laser ranging tests of the equivalence principle and PPN parameters; referenced in §5's PPN remark.]

Wagner, T. A., Schlamminger, S., Gundlach, J. H., & Adelberger, E. G. (2012). *Torsion-balance tests of the weak equivalence principle*. Classical and Quantum Gravity, **29**(18), 184002. [Modern Eötvös-style tests bounding violations of universal coupling; referenced in §10.2's Predicted falsification path.]

Will, C. M. (2014). *The confrontation between general relativity and experiment*. Living Reviews in Relativity, **17**, Article 4. [Comprehensive review of PPN parameter constraints; provides the empirical framework against which §5's PPN-β prediction is to be tested.]

Information-theoretic and thermodynamic foundations

Landauer, R. (1961). *Irreversibility and heat generation in the computing process*. IBM Journal of Research and Development, **5**(3), 183–191. [The Landauer bound on the thermodynamic cost of irreversible information erasure; underlies A2's "Landauer-like" content and the §9.5 conjectured A2–Definition 7 bridge.]

Bennett, C. H. (1982). *The thermodynamics of computation — a review*. International Journal of Theoretical Physics, **21**(12), 905–940. [Foundational review of the thermodynamics of irreversible computation; relevant to A2's thermodynamic content and the commitment-cost framework.]

Mathematical foundations

Ignatowski, W. v. (1910). *Einige allgemeine Bemerkungen über das Relativitätsprinzip*. Physikalische Zeitschrift, **11**, 972–976. [The von Ignatowski argument deriving Lorentz transformations from homogeneity, isotropy, and group structure together with frame-invariance of a maximum speed; used in Theorem 2 and Theorem A.]

Pal, P. B. (2003). *Nothing but relativity*. European Journal of Physics, **24**(3), 315–319. [Modern presentation of the Ignatowski-type derivation of Lorentz kinematics from minimal kinematic postulates; pedagogically relevant to Theorem 2.]

VERSF programme companion papers

The following are companion papers within the VERSF Theoretical Physics Programme (AIDA Institute; versf-eos.com) that this paper cites or depends upon. Full citations are available at the programme website; preprint and journal-publication details are pending and will be inserted at submission.

Taylor, K. *The K = 7 No-Go Theorem: Simplicial Closure and Boundary Architecture in the VERSF Substrate*. VERSF Theoretical Physics Programme. [Establishes the K = 7 wheel as the canonical closure cell; supplies §9.10's criteria (B) fold-interface compatibility and (C) primitive-eigenvalue normalisation.]

Taylor, K. *The Fold-Density Paper: Single-Fold Substrate and the Origin of Mass-Energy*. VERSF Theoretical Physics Programme. [Develops the single-fold substrate content that grounds (NPC) in §10.2's Theorem B; establishes that all matter reduces to fold configurations.]

Taylor, K. *The Commitment-Barrier Paper: Thermodynamics of Irreversible Record Formation*. VERSF Theoretical Physics Programme. [Companion to A2; develops the commitment-cost framework relevant to the §9.5 conjecture.]

Taylor, K. *The Law-Closure Paper: Observer-Protocols and the Origin of (OC)*. VERSF Theoretical Physics Programme. [Source of the observer-protocol primitives imported via (OC); supplies the axiomatic framework underlying Definition 6.1.]

Taylor, K. *Proto-Time and Emergent Lorentz Invariance*. VERSF Theoretical Physics Programme. [Develops the substrate-first construction of the temporal dimension referenced in §9.2 Step 5 and §10.3 Theorem C; addresses the Lorentzian-signature content that Theorem C explicitly does not derive.]

Taylor, K. *The Bit-Tick Ontology and TPB/BCB Foundations*. VERSF Theoretical Physics Programme. [Foundational programme content underlying the bit-versus-tick distinction in §9.4.6 and the Facthood = Invariant Comparability proposition in §2.1.]

Taylor, K. *Physics Without Time: Tick-Based Reformulation*. VERSF Theoretical Physics Programme. [Earlier programme work on the temporal-ontology questions addressed in this paper's Theorem C and Step 5 of Theorem 9.]

Programme website