Gravity from Tensorial Closure of Record Dynamics in the VERSF Framework

A leading-order uniqueness result for the gravitational response sector

---

Abstract for the General Reader

Modern physics describes gravity through the geometry of spacetime — a curved four-dimensional fabric whose shape is dictated by the matter and energy it contains. Einstein's equations make this precise, and a century of experiment has confirmed them. But the equations themselves are postulated, not derived: the rank-two metric tensor, its symmetric character, and the Einstein dynamics that govern it are simply assumed at the start.

The VERSF programme proposes a different starting point. It identifies physical reality not with a pre-existing geometry, but with the accumulation of *committed records* — irreversible "fact events" whose structure, when added up, gives rise to everything we ordinarily call the physical world. Time, space, fields, and forces are not primitive ingredients but emergent consequences of how committed records compose under a small set of structural constraints. Earlier papers in this programme have shown how quantum mechanics, the Born rule, and the basic kinematics of measurement can be derived from these foundations.

This paper addresses the gravitational sector. The question it asks is: given the closure constraints that make a record-based universe internally consistent, what is the *minimal* structure required to describe how committed records respond to one another at long distances — i.e., what does gravity have to look like? The result is that:

- The response must be encoded in a symmetric rank-two tensor field (the same mathematical object as Einstein's metric), not a scalar, vector, or any higher structure. This is forced by representation theory: only this particular kind of object can couple in the right way to all the ingredients of the energy-momentum source.
- The conservation law that this tensor must satisfy takes a specific, unique form at leading order.
- A geometric metric — and hence the curvature picture of gravity — emerges as a consequence of how committed records can be consistently transported across infinitesimally separated points.
- Einstein's field equations themselves emerge as the unique low-energy description of this geometric sector.

The paper is honest about the inputs it requires beyond the closure constraints. Five supplementary corpus inputs are required — the existence of a smooth four-dimensional manifold, its Lorentzian signature, an effective-field-theory power-counting structure, a substrate-level ordering of commitment events, and a substrate-level commutativity property of

pointwise commitments. Each is motivated as the continuum-limit expression of a feature already implicit in record-based physics, rather than a free assumption; their full derivation is the proper subject of subsequent papers. Within these inputs, gravity is not an independent force but the unique tensorial closure of record dynamics consistent with a universe of irreversible facts.

For the working physicist, the paper offers four sharp falsification criteria, including specific signatures in gravitational-wave polarisation, dispersion, and equivalence-principle tests for spinning bodies — all of which would, if observed, force revision of identified components of the construction.

Table of Contents

Technical Abstract

The VERSF programme identifies physical reality with the accumulation of irreversible commitment events forming a structured record. Scalar closure of the record density yields a unique conservation law and a Klein–Gordon-type equation of motion, but is structurally insufficient to encode gravitational phenomena, which require directional and relational response between independent axes.

In this paper we derive the minimal tensorial extension of record dynamics required for gravitational response on the *already-emerged* 4-manifold M. Under six closure constraints —

finite distinguishability, irreversibility, locality, additivity, observer invariance, and minimality (formalised as a well-defined extremisation principle), supplemented by five corpus inputs (smooth manifold M, Lorentzian signature, EFT power-counting, substrate-level commitment ordering, and pointwise commitment commutativity — labelled (E1)–(E5) in §12.2) which together serve as the minimal bridge from the substrate-level record ontology to the post-emergence smooth theory, we establish the following structural results:

1. **Rank selection (Theorem 1).** Classification of admissible response sectors under (C1)–(C6) reduces to a classification of irreducible representations of SO(3,1) admitting a universal linear coupling to the committed energy-momentum source. By Schur's lemma — at bilinear and zero-derivative order, with both restrictions reduced to (C6) — the unique minimal admissible irrep is the symmetric rank-2 representation $(1,1) \oplus (0,0)$.
2. **Symmetry derivation.** The symmetry of the response sector descends from Theorem 1; it is not postulated.
3. **Linear leading-order uniqueness (Theorem 2).** Within the linear, local, leading-derivative-order (zero-derivative-in-Φ) ansatz, the leading-order admissible record current is uniquely

   $$\mathcal{C}^{\mu\nu} = a\,\Phi^{\mu\nu} + b\,g^{\mu\nu}\,\Phi\,,$$

   with $\Phi \equiv g^{\mu\nu}\Phi_{\mu\nu}$ the metric trace of the fundamental tensorial commitment-density field. The current $\mathcal{C}^{\mu\nu}$ is a *linear projection* of $\Phi_{\mu\nu}$ (analogous to the trace-reversed metric perturbation of linearised GR), distinct from the stress-energy tensor $T^{(\Phi)}_{\mu\nu}$ which is quadratic in Φ.

4. **Metric emergence (Theorem 3).** Closure-consistent parallel transport of committed structure, combined with the supplementary input (E5) of pointwise commitment commutativity, determines a unique torsion-free linear connection on the substrate. Non-degeneracy of $\Phi_{\mu\nu}$ (forced by (C4) universality) identifies $\Phi_{\mu\nu}$ as a metric, and the fundamental theorem of Riemannian geometry then identifies the connection as its Levi-Civita connection. The metric $g_{\mu\nu} = \Phi_{\mu\nu}/\lambda_\star$ is unique up to the normalisation set by the finite-distinguishability scale $\ell_\star$.
5. **Einstein structure.** Lovelock's theorem applied to the emergent geometric sector selects the Einstein form

   $$G_{\mu\nu} + \Lambda\,g_{\mu\nu} = \kappa\,T^{(\Phi)}_{\mu\nu}$$

   as the unique low-derivative covariant closure, sourced by the stress-energy tensor $T^{(\Phi)}_{\mu\nu}$ (quadratic in Φ), not by $\mathcal{C}^{\mu\nu}$.

We exhibit the master action $\mathcal{S}[\Phi_{\mu\nu}, g_{\mu\nu}, \psi]$ in its post-emergence effective form (presupposing Theorem 3), with proper Fierz–Pauli kinetic terms, and fix the dimensional bookkeeping by canonical normalisation: $[\Phi_{\mu\nu}] = M^1$, $[\mathcal{C}^{\mu\nu}] = M^1$, $[T^{(\Phi)}_{\mu\nu}] = M^4$. *The corpus identity* $TPB[\Phi] \sim |\Phi|^2$ *is realised at the schematic level via the same* $M_\star$ bridge factor that mediates matter coupling — consistent with minimality, since a different scale would introduce an additional physical scale not present in (C1). We extract four sharp falsification criteria, including detection of antisymmetric rank-2 phenomenology in gravitational-wave polarisation modes (F1), gravitational-wave dispersion at LIGO–Virgo–KAGRA frequencies (F2), low-density inverse-square deviations off the predicted parametric form (F3), and torsion in spinning-body equivalence-principle tests (F4). We close with an explicit logical dependency structure separating axioms, supplementary inputs, derived results, conditional claims, and ordering relations.

---

1. Introduction

The standard formulation of gravity treats the rank-2 symmetric metric tensor and its Einstein dynamics as primitive postulates. The VERSF programme inverts this hierarchy: geometric, dynamical, and field-theoretic structure are taken to *emerge* from the accumulation of irreversible commitment events under a small set of admissibility constraints — finite distinguishability, irreversible commitment, locality, additivity, observer invariance, and minimality.

Earlier papers in the corpus established a conserved record current at scalar order, $\partial_\mu \mathcal{C}^\mu = s_c$; a unique scalar dynamical equation under minimal closure, $(\Box + m^2)\,\Phi = \sigma$; and a derivation of quantum kinematics, the Born rule, and unitary evolution from finite distinguishability and irreversible commitment.

Scalar closure cannot account for gravity. Gravitational phenomena exhibit anisotropic response, tidal structure, and curvature — features that demand encoding *relations between directions*, not just magnitudes assigned to points. The present paper therefore asks:

Given the VERSF closure constraints, what is the minimal tensorial extension of record dynamics, and how uniquely is its leading-order form determined?

The contribution is fivefold. We formalise (C6) as an extremisation principle on the space of admissible response functionals, applicable in §4 without forward reference to downstream conservation identities. We reframe rank selection as a classification problem on irreducible representations of the Lorentz group, replacing case-by-case dismissal with a Schur-type uniqueness argument. We derive — rather than postulate — both the symmetric character of the response sector and the existence of a metric: the former as a consequence of the rank-2 selection theorem, the latter as Theorem 3 via closure-consistent transport given a supplementary input (E5) on substrate-level pointwise commitment commutativity, which is explicitly listed as a corpus dependency rather than smuggled in. We write the master action explicitly, fixing the dimensional bookkeeping required by the corpus identity $\mathrm{TPB}[\Phi] \sim |\Phi|^2$. And we close with a logical dependency table making explicit what is axiomatic, what is supplementary, what is derived, and what is conditional.

We are deliberately scope-honest throughout. The uniqueness theorems below are leading-order results within explicitly stated ansatz classes.

2. Why Scalar Closure Is Insufficient

Scalar closure assigns a single real number $\Phi(x)$ to each spacetime event. The associated conservation law

$$\partial_\mu \mathcal{C}^\mu = s_c ,$$

and the leading scalar dynamical equation

$$(\square + m^2)\, \Phi = \sigma ,$$

describe accumulation and propagation of committed structure but have three structural defects with respect to gravitational phenomena.

No directional dependence. A scalar field has no preferred axis at any point, so it cannot encode anisotropic response. Tidal forces — which by definition distinguish between directions transverse and parallel to the source — cannot be sourced by Φ alone.

No relational encoding. Curvature is intrinsically pairwise: it measures how parallel transport of a vector around an infinitesimal loop fails to return to itself, an object requiring two independent directions to define. A scalar cannot encode such a pairwise relation.

Insufficient coupling structure. A scalar response field can couple universally only to the trace of the committed energy-momentum tensor; commitments with vanishing trace are gravitationally invisible to it. This is incompatible with universality, which §4 will show is forced by additivity (C4).

We conclude that scalar closure exhausts the *magnitude* sector of record dynamics but cannot describe its *relational* sector.

---

3. Closure Constraints

We restate the constraints, with (C6) sharpened.

(C1) Finite distinguishability. All physical quantities are finitely resolvable; the substrate admits a finite distinguishability scale $\ell_\star$.

(C2) Irreversible commitment. Dynamics monotonically accumulate committed structure; the commitment-density functional $\mathcal{S}[\Phi]$ is non-decreasing along the substrate-level ordering of commitment events from which the emergent temporal coordinate of the manifold M derives.

The notion of "history" here therefore refers to the substrate-level ordering, not to a pre-existing temporal coordinate.

(C3) Locality. Response at x depends only on commitment structure within an infinitesimal neighbourhood of x; the response functional admits a derivative expansion.

(C4) Additivity. Commitments from causally independent regions contribute additively to the record. Equivalently, the response functional is linear at leading order in the source decomposition, and the source of the response field is the *universal* committed energy-momentum tensor $T_{\mu\nu}$ (no commitment density is excluded).

We fix terminology: a coupling between a response field F and the source $T_{\mu\nu}$ is called **universal** if and only if it is a non-trivial Lorentz-invariant local interaction that depends on *all* components of $T_{\mu\nu}$ — both its trace $g^{\mu\nu} T_{\mu\nu}$ and its trace-free part $T_{\mu\nu} - (\frac{1}{4}) g_{\mu\nu} g^{\alpha\beta} T_{\alpha\beta}$. A coupling that depends only on the trace, only on the trace-free part, or only on a proper subset of the source's irreducible components is *non-universal* and violates (C4). This notion of universality is the structural pivot of the rank-selection argument: every elimination in §4.3 is driven by the failure of a candidate irrep to support universal coupling in the sense defined here.

(C4′) Configuration-space inclusion. Theorem 3's non-degeneracy step (§9.2) requires a strengthening of (C4) beyond the bare irrep-level statement. Specifically, the matter sector includes commitments structured along *every* tangent direction — i.e., $T_{\mu\nu}$ configurations of the form $v_\mu v_\nu$ for any $v \in T_x M$ are admissible, not merely those restricted to a sub-cone of the symmetric rank-2 cone. We label this strengthening (C4′) and treat it as a sub-input of (C4): the universality of (C4) at the irrep level does not by itself imply that all configurations within the admissible representation are realised, and (C4′) supplies the additional configuration-space content. The non-degeneracy argument of §9.2 invokes (C4′), not (C4) alone; (C4′) is flagged in §12 as an explicit dependency of Theorem 3.

(C5) Observer invariance. No preferred frame, basis, or labelling enters the response; all admissible response functionals are diffeomorphism-covariant on the emergent 4-manifold M (whose smooth structure and Lorentzian signature are corpus inputs, not derived here; see end of this section). In Minkowski tangent spaces this reduces to global SO(3,1) covariance.

(C6) Minimality (formalised). Among response functionals satisfying (C1)–(C5), admissible response is the extremum of the *operator-content functional*

$\mathcal{N}[F] = \sum_{\mathcal{O}}$ (number of independent local operators in F at order $\leq$ N),

at each derivative order N consistent with reproducing the universal coupling structure to that order. Equivalently: no operator may appear in F that is independent of the irreducible Lorentz-component structure required to support universal coupling (per the §4.2 definition) to $T_{\mu\nu}$.

This formulation is intentional. (C6) refers only to the universal coupling structure and its irreducible Lorentz components, both of which are intrinsic to (C3)–(C5) and require no input from §8 or §9. (C6) is consequently a pure operator-content minimality condition, applicable in §4 without forward reference to downstream constructions. The conservation-identity formulation that one might naturally write — referring to closure of $\nabla_\nu \mathcal{C}^{\mu\nu} = \mathcal{S}^\mu$ — is recovered downstream as a *consistency check* on the construction (§8.4), not as the definition of (C6). Defining (C6) via the conservation identity would forward-reference §8 (which presupposes §9), creating a cyclic dependency when (C6) is invoked in §4 to derive Theorem 1; the upstream operator-content formulation given here avoids that cycle.

This converts "minimality" from a heuristic to a constraint with definite content, and locates that content cleanly in the upstream sector of the dependency graph.

Corpus imports presupposed by (C1)–(C6)

The constraints (C1)–(C6) are the inputs *internal* to this paper. They are layered on top of corpus-level results not re-derived here. In particular:

- the emergence of a smooth 4-manifold M from substrate dynamics — the precondition for taking derivatives, defining smoothness, and expanding response functionals as derivative series (presupposed by (C3));
- the Lorentzian signature of M's tangent-space structure — the precondition for replacing diffeomorphism covariance with SO(3,1) covariance in Minkowski tangent spaces (presupposed by (C5));
- the substrate-level ordering of commitment events from which the emergent temporal coordinate derives — the precondition for "monotonic accumulation" being a well-defined statement (presupposed by (C2));
- pointwise commutativity of infinitesimal commitment composition — required for the metric-emergence theorem of §9 (introduced and discussed in §9.1).

These are formalised as (E1), (E2), (E4), (E5) in §12.2, where each is motivated as the continuum-limit expression of a structural feature already implicit in record-based physics — i.e., as a minimal bridge assumption from the substrate ontology to a smooth post-emergence theory, not as a free assumption. An additional supplementary input — EFT power-counting in $(\Phi/\Phi_\star)$ — is introduced in §6 where it is first invoked, and listed as (E3) alongside the others in §12.2.

The construction below is therefore a theory of *dynamics on the emergent manifold*, not of its emergence. A fully substrate-up derivation (deriving the manifold and the dynamics on it in a single construction) is the subject of Paper 3.

---

4. Rank Selection as Classification of Lorentz Irreps

We reformulate rank selection as a representation-theoretic classification problem, replacing case-by-case framing with a single structural result.

4.1 The classification problem

By (C5), the response field F transforms in some finite-dimensional representation R_F of SO(3,1) (equivalently, of its double cover $SL(2,\mathbb{C})$, so we label irreps as (j_L, j_R) with $j_L, j_R \in \frac{1}{2}\mathbb{Z}_{\geq 0}$). By (C4), the source of F is the universal committed energy-momentum tensor $T_{\mu\nu}$, which transforms in the representation $R_T = (1,1) \oplus (0,0)$. By (C3) and (C5), the leading interaction term is a local Lorentz scalar bilinear in F and $T_{\mu\nu}$.

The question is: which irreps R_F admit a non-trivial Lorentz-invariant local coupling to R_T?

4.2 The intertwiner constraint

A Lorentz-invariant local coupling $\mathcal{L}_{int}(F, T) \in \mathbb{R}$ exists if and only if $R_F \otimes R_T$ contains the trivial representation (0,0). For real-valued fields this is equivalent (by self-duality of finite-dimensional SO(3,1) reps under the conjugation $(j_L, j_R) \to (j_R, j_L)$ for reps with real characters) to

$$R_F \supseteq R_T = (1,1) \oplus (0,0).$$

The minimal such R_F is R_T itself. This is the symmetric rank-2 representation: a symmetric rank-2 tensor decomposes as (1,1) [traceless symmetric] ⊕ (0,0) [trace].

Lemma (Schur-type, leading-order). *Under (C3)–(C5), restricted to bilinear leading-order interactions in F and T, an admissible response field F admits a non-trivial Lorentz-invariant universal coupling to $T_{\mu\nu}$ only if R_F contains both (1,1) and (0,0) as sub-representations — equivalently, $R_F \supseteq (1,1) \oplus (0,0)$.*

By the §3 definition of universality, this requires R_F to pair non-trivially with both the (1,1) and (0,0) sectors of T (the trace-free and trace pieces respectively). The lemma's force rests entirely on this universality requirement: a non-universal coupling — e.g., a scalar φ coupling only to the trace via $\varphi \cdot g^{\mu\nu} T_{\mu\nu}$ — exists as a Lorentz scalar bilinear interaction and is not excluded by Schur's lemma alone. The exclusion of such non-universal candidates is content of (C4), not of Schur.

Scope of the lemma. Two qualifications are required to make the scope honest:

(a) Bilinear restriction. The argument assumes the leading interaction is *bilinear* — linear in F and linear in T. Higher-multilinear interactions, such as a coupling $F^{\mu\rho} F_{\rho}{}^{;\nu} T_{\mu\nu}$ that is quadratic in F, admit additional intertwiners that the bilinear analysis does not detect (e.g., a vector field A_μ admits a coupling $A_\mu A_\nu T^{\mu\nu}$, which is a Lorentz scalar). Such higher-multilinear interactions are excluded at leading order by (C6) minimality: they introduce operator content beyond the irreducible Lorentz-component structure required for universal coupling at zero-derivative bilinear order. The bilinear restriction is therefore a *consequence* of (C6) at leading order, not an independent assumption. At subleading orders, multilinear interactions become available and contribute to the suppressed corrections of §11.1.

(b) Zero-derivative restriction. The argument applies at zero derivative order in F. With one or more derivatives included, the index-bearing object $\partial_{\mu_1} \cdots \partial_{\mu_k} F$ transforms in a tensor product of representations that may contain $(1,1) \oplus (0,0)$ even when R_F itself does not. We address the most consequential case — vector fields — explicitly in §4.3 below. As with the bilinear restriction, derivative couplings are excluded at leading order by (C6) and suppressed by inverse powers of $M_\star$ at subleading order.

The reduction from containment to equality is supplied by (C6). Suppose $R_F = (1,1) \oplus (0,0) \oplus R_{extra}$ for some non-trivial R_{extra}. The components of F transforming in R_{extra} are independent of the irreducible Lorentz-component structure required for universal coupling to $T_{\mu\nu}$, since by definition the universal coupling is exhausted by the $(1,1) \oplus (0,0)$ sector matching $T_{\mu\nu}$. Per (C6) — in its operator-content-minimality formulation per §3, which does not invoke the conservation identity — R_{extra} carries operator content not required by the universal coupling structure. By the extremisation principle of (C6), R_{extra} must be excluded at leading order.

Corollary. *Under (C1)–(C6), $R_F = (1,1) \oplus (0,0)$ exactly — equality, not merely containment, at leading order in (C6)'s ansatz class.*

4.3 Elimination of competing irreps

Applying the lemma:

- **Scalar (0,0).** $R_F = (0,0)$ does not contain (1,1). It admits coupling only to the trace sector of $T_{\mu\nu}$, leaving the traceless symmetric part decoupled. Violates (C4).
- **Vector (½, ½).** $R_F = (½, ½)$ contains neither (1,1) nor (0,0). The intertwiner with R_T is zero at zero-derivative order — a bilinear coupling $A_\mu T^{\mu\nu} \cdots$ requires the response to contract with two indices of T, which a vector cannot do. *Caveat at one-derivative order:* with one derivative included, $\partial_\mu A_\nu$ transforms as $(½,½) \otimes (½,½) = (1,1) \oplus (1,0) \oplus (0,1) \oplus (0,0)$, which *does* contain $(1,1) \oplus (0,0)$. A coupling $\partial_\mu A_\nu T^{\mu\nu}$ is therefore a valid Lorentz scalar at one-derivative order. However, this coupling is on-shell trivial for a conserved source: integration by parts gives

$$\int \partial_\mu A_\nu T^{\mu\nu} = \int \partial_\mu(A_\nu T^{\mu\nu}) - \int A_\nu \partial_\mu T^{\mu\nu} ,$$

where the first term is a boundary contribution and the second vanishes for $\partial_\mu T^{\mu\nu} = 0$ (which holds in the relevant matter sector by the conservation of the symmetric stress-energy tensor; the argument extends straightforwardly to $\nabla_\mu T^{\mu\nu} = 0$ in curved space, where the boundary term vanishes for asymptotically-vanishing fields). The one-derivative vector coupling is therefore not subleading but identically zero on-shell, providing a structural exclusion independent of (C6) minimality. Thus vector-mediated coupling is not forbidden in principle, but vanishes on-shell at one-derivative order and is excluded at leading order by (C6). Violates (C4) at leading order.

- **Antisymmetric rank-2 (1,0) ⊕ (0,1).** R_F contains neither (1,1) nor (0,0). Antisymmetric rank-2 fields couple universally to *antisymmetric* sources ($B_{\mu\nu}$

$H^{\mu\nu}$ with $H_{\mu\nu}$ antisymmetric), not to the symmetric universal source $T_{\mu\nu}$. The intertwiner with R_T is zero. Violates (C4).

- **Symmetric rank-2 (1,1) ⊕ (0,0).** $R_F = R_T$ exactly. Admits universal coupling $\Phi^{\mu\nu} T_{\mu\nu}$. ✓
- **Higher rank (j_L, j_R) with $j_L + j_R \geq 3/2$ not equal to (1,1) ⊕ (0,0).** Either does not contain (1,1) ⊕ (0,0) at all (excluded by the lemma), or contains it as a proper sub-representation alongside additional irreducible components carrying independent operator content (excluded by (C6) minimality at leading order).

The vector rejection in particular is no longer a sign-of-static-interaction argument — which is vulnerable to coupling-convention objections — but a representation-theoretic impossibility: the intertwiner literally does not exist.

4.4 Rank selection theorem

Theorem 1 (Rank Selection). *Under (C1)–(C6), the unique minimal admissible response sector for universal coupling to the committed energy-momentum source is the symmetric rank-2 representation (1,1) ⊕ (0,0) of SO(3,1) — equivalently, a symmetric rank-2 tensor field $\Phi_{\mu\nu}$.*

The rank-2 structure is therefore not postulated; it is the unique admissible irrep selected by the closure constraints under the Schur-type criterion of §4.2.

4.5 Remark on antisymmetric rank-2

The antisymmetric sector is not forbidden as a physical field — it is consistent and may carry independent content (e.g., torsional or Kalb–Ramond-like structure). What §4.3 establishes is that it cannot serve as the *gravitational* response, since gravity is precisely the response coupled to the universal symmetric source. The distinction is structural: gravitational response is the (1,1) ⊕ (0,0) sector by the §3 definition of universal coupling to $T_{\mu\nu}$, not by an additional postulate.

---

5. The Tensorial Commitment-Density Field

We define the fundamental tensorial object and its relationship to the scalar Φ of Paper 1.

Let $\Phi_{\mu\nu}(x)$ denote the **tensorial commitment-density field** — a symmetric rank-2 covariant field whose components measure the directionally-resolved density of committed structure at x. The scalar field of the earlier scalar-closure paper is recovered as the metric trace,

$$\Phi \equiv g^{\mu\nu} \Phi_{\mu\nu} ,$$

so the prior scalar dynamics is the trace projection of the tensorial sector. This identification is the only one consistent with three requirements: that the scalar closure of Paper 1 be the trace sector of the present tensorial closure; that $\Phi_{\mu\nu}$ reduce to the scalar regime when the directional structure of the source is isotropic; and that the dimensional bookkeeping fixed in §10 close consistently.

In particular, $\Phi_{\mu\nu}$ is *not* defined as $\nabla_\mu \nabla_\nu \Phi$. That would make the tensorial sector derivative of the scalar one, whereas Theorem 1 establishes that the scalar sector is the *trace projection* of an independent tensorial sector. The relationship is that of trace to full tensor, not that of derived to fundamental.

---

6. Linear Uniqueness of the Leading-Order Record Current

We construct the most general local covariant rank-2 functional of $\Phi_{\mu\nu}$ satisfying (C1)–(C6) at the linear, leading-derivative-order (zero-derivative-in-Φ) level. Linearity follows from (C4) at leading order. Locality (C3) restricts to operators of finite derivative order. Observer invariance (C5) restricts to diffeomorphism scalars. The available structures at zero derivatives in $\Phi_{\mu\nu}$ are exhausted by

$$\Phi^{\mu\nu} , \; g^{\mu\nu} \Phi ,$$

and any other zero-derivative covariant rank-2 object built from $\Phi_{\alpha\beta}$ and $g_{\alpha\beta}$ reduces by linearity to a linear combination of these two. Therefore the most general admissible leading-order current is

$$\mathcal{C}^{\mu\nu} = a\, \Phi^{\mu\nu} + b\, g^{\mu\nu} \Phi ,$$

with dimensional coefficients a, b fixed by the master action (§8) and dimensional bookkeeping (§10).

Theorem 2 (Leading-Order Linear Uniqueness). *Within the class of local, covariant, zero-derivative-in-Φ functionals linear in $\Phi_{\mu\nu}$, with higher-derivative corrections suppressed by inverse powers of $M_\bigstar$ per (E3), the leading-order record current is uniquely*

$$\mathcal{C}^{\mu\nu} = a\, \Phi^{\mu\nu} + b\, g^{\mu\nu} \Phi.$$

The "leading-order" qualifier is the parametric leading-order in $(\Phi/\Phi_\bigstar)$ and $(\partial/M_\bigstar)$, which presupposes (E3); the linear ansatz itself is an immediate consequence of (C4) at zero order. At non-zero derivative order in Φ, additional admissible structures appear — $\nabla^\mu \nabla^\nu \Phi$, $\nabla^2 \Phi^{\mu\nu}$, $\nabla^\mu \nabla_\alpha \Phi^{\alpha\nu}$, $g^{\mu\nu} \nabla^2 \Phi$, $g^{\mu\nu} \nabla_\alpha \nabla_\beta \Phi^{\alpha\beta}$ *— all of which are suppressed by $1/M_\bigstar^2$* per (E3) and contribute to the parametrically subleading corrections of §11.

We state the scope precisely. This is uniqueness within the linear ansatz. Nonlinear operators of the schematic forms $\Phi^{\mu\alpha} \Phi_\alpha{}^{;\nu}$, $\Phi^2 g^{\mu\nu}$, and $(\nabla\Phi)^2 g^{\mu\nu}$ are not excluded by (C1)–(C6) but are suppressed by powers of $(\Phi/\Phi_\bigstar)$, where $\Phi_\bigstar$ is the commitment density at the finite-distinguishability scale $\ell_\star$ from (C1). At commitment densities $\Phi \ll \Phi_\star$ — which includes all currently observed gravitational regimes — the linear ansatz is parametrically justified. The breakdown regime is discussed in §11.

Status of the suppression-scale claim. The statement that nonlinear operators are suppressed by powers of $(\Phi/\Phi_\star)$ with O(1) Wilson coefficients is the standard naive-dimensional-analysis (NDA) organisation of an effective Lagrangian as a power series in field amplitudes over a UV scale. Within the present construction, the finite-distinguishability scale $\ell_\star$ from (C1) supplies the dimensional scale $\Phi_\star$, but the *power-counting structure* — that operator coefficients of mass dimension n are bounded by $M_\star^{-n}$ with O(1) Wilson coefficients — is an additional dynamical assumption beyond (C1)–(C6). It is the analogue of NDA in standard EFT and would, in a full UV completion, follow from the matching procedure to the substrate-level dynamics. For the present paper, EFT power-counting is taken as a *supplementary corpus input*, listed as (E3) in §12.2. The substrate-level derivation of (E3) is the subject of Paper 3.

---

7. Symmetry Derivation

The symmetry of $\mathcal{C}^{\mu\nu}$ under $\mu \leftrightarrow \nu$ is not a free choice. It descends from the symmetry of $\Phi_{\mu\nu}$ (Theorem 1) and the absence of an antisymmetric covariant rank-2 structure linear in $\Phi_{\mu\nu}$ and $g_{\mu\nu}$. Explicitly, any candidate antisymmetric contribution would be of the form

$$\alpha\,(\Phi^{\mu\nu} - \Phi^{\nu\mu}) + \beta\,(g^{\mu\nu} - g^{\nu\mu})\,\Phi\,,$$

both of which vanish identically given the symmetry of $\Phi_{\mu\nu}$ and $g_{\mu\nu}$. Therefore

$$\mathcal{C}^{\mu\nu} = \mathcal{C}^{\nu\mu}$$

automatically, and no additional symmetrisation postulate is required.

---

8. The Master Action and Conservation Identity

8.1 Logical placement and the background/perturbation split

The master action below is the *post-emergence effective description*, valid in the regime where the metric $g_{\mu\nu}$ has been derived per Theorem 3 of §9 and treated as a fixed background. The action functional therefore presupposes the geometric structure that §9 establishes from closure-consistent transport. A pre-emergence formulation — in which only the fundamental tensorial commitment-density field appears and the metric structure has not yet been constructed — would be expressed in terms of substrate-level operators on the bare field, with no $\sqrt{-g}$ measure or $g^{\mu\nu}$ contractions. We do not write that pre-emergence action here; it would require a treatment of substrate-level dynamics that produces Theorem 3, which is the subject of Paper 3 in the roadmap.

§8 logically follows §9, even though it is presented before it for expository clarity. The dependency table in §12 reflects this.

Background/perturbation split. A subtlety arises from §9's identification $\Phi_{\mu\nu} = \lambda_\star(x)\, g_{\mu\nu}(x)$: the fundamental tensorial commitment-density field *is* the metric (up to local scale). Writing an action for "$\Phi_{\mu\nu}$ on background $g_{\mu\nu}$" without further specification would amount to writing an action for a field on its own background. We resolve this by the standard linearised-gravity split. The fundamental field of §9, denoted $\bar{\Phi}_{\mu\nu}$, decomposes into a chosen background plus a propagating perturbation:

$$\bar{\Phi}_{\mu\nu}(x) = \lambda_\star(x)\, \tilde{g}_{\mu\nu}(x) + h_{\mu\nu}(x) ,$$

where $\tilde{g}_{\mu\nu}$ is the chosen background metric (equivalently, $\bar{\Phi}/\lambda_\star$ in the absence of fluctuations) and $h_{\mu\nu}$ is the propagating perturbation. The action of §8.2 governs $h_{\mu\nu}$ on the background $\tilde{g}_{\mu\nu}$. For brevity we write $\Phi_{\mu\nu}$ for $h_{\mu\nu}$ throughout §8 — the symbol denotes the perturbation, not the fundamental field of §9. Where the distinction matters (e.g., in the relation between $\mathcal{C}^{\mu\nu}$ and the source structure), we make it explicit.

This is the standard treatment of a propagating tensor on a background that is itself a coarse-grained projection of the same field — the analogue of writing a graviton action $h_{\mu\nu}$ on a fixed background metric $\tilde{g}_{\mu\nu}$ in linearised general relativity, where $\tilde{g}$ is the background piece of the full metric $g = \tilde{g} + h$.

8.2 The post-emergence action

Define

$$\mathcal{S}[\Phi_{\mu\nu}, g_{\mu\nu}, \psi] = \int d^4x \sqrt{-g} \left[-(1/2)\, \nabla_\lambda \Phi^{\mu\nu} \nabla^\lambda \Phi_{\mu\nu} + (\kappa_1/2)\, (\nabla_\mu \Phi^{\mu\nu})(\nabla^\lambda \Phi_{\lambda\nu}) + (\kappa_2/2)\, (\nabla_\mu \Phi)(\nabla^\mu \Phi) - (m^2/2)\, \Phi^{\mu\nu} \Phi_{\mu\nu} + \mathcal{L}_{int}(\Phi_{\mu\nu}, \psi; g_{\mu\nu}) \right] ,$$

where $\Phi_{\mu\nu}$ is the perturbation field $h_{\mu\nu}$ of §8.1 (the symbol is reused for compactness; see §8.1 for the relation to the fundamental $\bar{\Phi}_{\mu\nu}$ of §9), the kinetic terms have the standard Fierz–Pauli-like structure for a propagating symmetric rank-2 field on the emergent background $g_{\mu\nu}$ ($\equiv \tilde{g}_{\mu\nu}$ of §8.1), the κ_1 and κ_2 coefficients fix the gauge structure of the linearised theory, m is a mass scale (which may vanish in the massless-graviton limit), and $\mathcal{L}_{int}$ couples $\Phi_{\mu\nu}$ to the matter sector ψ. The kinetic term is the structure that fixes canonical normalisation in §10.

8.3 The record current as a linear projection

A central conceptual point: the record current $\mathcal{C}^{\mu\nu}$ of Theorem 2 is **not** the stress-energy tensor of the Φ field. It is a *linear projection* of $\Phi_{\mu\nu}$, defined directly as

$\mathcal{C}^{\mu\nu} \equiv a\, \Phi^{\mu\nu} + b\, g^{\mu\nu}\, \Phi$,

with a, b dimensionless coefficients fixed by the matter-coupling structure. The form of $\mathcal{C}^{\mu\nu}$ is analogous to the trace-reversed metric perturbation $\bar{h}_{\mu\nu} = h_{\mu\nu} - (½)\, g_{\mu\nu}\, h$ of linearised general relativity, where the coefficients (a, b) = (1, −½) are fixed by the requirement that $\bar{h}_{\mu\nu}$ satisfy $\partial^\nu \bar{h}_{\mu\nu} = 0$ in the Lorenz gauge for a free field. The same logic applies here: the values of (a, b) are fixed by the requirement that $\mathcal{C}^{\mu\nu}$ satisfy a covariant divergence identity sourced by the matter sector.

8.4 The conservation identity

The conservation identity

$\nabla_\nu\, \mathcal{C}^{\mu\nu} = \mathcal{S}^\mu$,

follows from the equations of motion for $\Phi_{\mu\nu}$ (obtained by variation of $\mathcal{S}$ with respect to $\Phi_{\mu\nu}$), combined with the diffeomorphism invariance of $\mathcal{S}$. By Noether's second theorem, diffeomorphism invariance of the matter sector $\mathcal{L}_{int}$ implies that the matter-sector source

$\mathcal{S}^\mu \equiv$ functional of the matter source coupling to $\Phi_{\mu\nu}$

satisfies $\nabla_\nu\, \mathcal{C}^{\mu\nu} = \mathcal{S}^\mu$ on shell. The identity is therefore *automatic* given (C5), not an independent postulate.

8.5 The stress-energy tensor (distinct from $\mathcal{C}^{\mu\nu}$)

The stress-energy tensor of the Φ field is a distinct object,

$T^{(\Phi)}_{\mu\nu} \equiv -(2/\sqrt{-g})\, \delta\mathcal{S} / \delta g^{\mu\nu}$,

which is *quadratic* in Φ for a quadratic action, and has dimension $[T^{(\Phi)}_{\mu\nu}] = M^4$ *(consistent with energy density). It is* $T^{(\Phi)}_{\mu\nu}$ — not $\mathcal{C}^{\mu\nu}$ — that sources the Einstein equation derived in §9.3. The distinction is essential:

- $\mathcal{C}^{\mu\nu}$ is *linear* in Φ, has dimension M^1 (per §10), and satisfies the conservation identity.
- $T^{(\Phi)}_{\mu\nu}$ is *quadratic* in Φ, has dimension M^4, and sources geometric curvature.

It is essential to keep these two objects distinct. The full coefficient algebra relating (a, b, κ_1, κ_2, m) to the matter-sector Wilson coefficients is collected in Paper 4.

---

9. Metric Emergence

This section states metric emergence as a theorem rather than a definition.

9.1 From committed structure to a connection

Consider parallel transport of a committed reference frame between two infinitesimally separated events x and x + dx. By (C4), the transport map $P(x \to x + dx): T_x M \to T_{x+dx} M$ is linear; by (C3), it depends only on local commitment structure at x and dx. It therefore admits a derivative expansion

$$P^\mu{}_{\nu}(x \to x + dx) = \delta^\mu{}_{\nu} - \Gamma^\mu{}_{\nu\lambda}(x)\, dx^\lambda + O(dx^2)\ ,$$

defining a linear connection $\Gamma^\mu{}_{\nu\lambda}$ on M. The expansion is well-defined for any linear connection — torsionful or torsion-free — so the existence of P does not by itself constrain torsion. The choice between a torsionful Cartan-type connection and a torsion-free Levi-Civita connection is a separate question, addressed below.

Torsion-freeness. The connection a priori admits both a symmetric component $\Gamma^\mu{}_{(\nu\lambda)}$ and an antisymmetric component $T^\mu{}_{\nu\lambda} \equiv \Gamma^\mu{}_{[\nu\lambda]}$ (the torsion). A naive argument from infinitesimal loop closure is circular, since torsion is defined as the failure of infinitesimal parallelograms to close; we therefore need a structural input distinct from loop-closure itself.

We now give a derivation that is honest about its inputs. The non-trivial geometric question is what distinguishes the *commutator of infinitesimal commitments at a point* (whose vanishing is torsion-freeness) from the *holonomy of parallel transport around a region* (whose non-vanishing is curvature). The former is excluded; the latter is not. This separation requires a structural input about the substrate-level commitment algebra that is not contained in (C1)–(C5).

The required input. Two infinitesimal commitment events dC_1 and dC_2 at a single substrate point x along directions e_μ and e_ν must compose commutatively as elements of the local commitment algebra at x:

$$dC_1 \circ dC_2 = dC_2 \circ dC_1\ .$$

This pointwise commutativity is *not* a consequence of (C1)–(C5). One might attempt to derive it from (C2) (the record encodes only what is committed) plus (C3) (response depends only on local structure). That derivation is suggestive but presupposes the identification of "what is committed at x" with the additive accumulated density at x — which is precisely what pointwise commutativity of the substrate algebra asserts. In a non-commutative substrate (analogue: non-commuting quantum measurements at a point), the record could in principle distinguish $dC_1 \circ dC_2$ from $dC_2 \circ dC_1$ at first order.

We therefore promote pointwise commutativity to a supplementary corpus input:

(E5) Pointwise commitment commutativity. Infinitesimal commitments at a single substrate point compose commutatively as elements of the local commitment algebra. Equivalently, the substrate-level commitment composition is abelian at first order in the resolution scale $\ell_\star$.

This is added to the supplementary-inputs list in §12.2. Its derivation from substrate dynamics is the subject of Paper 3.

Derivation of torsion-freeness from (E5). Given (E5), commutativity of pointwise commitment composition translates, by linearity of the transport map P (per (C4)) and locality of P (per (C3)), into commutativity of infinitesimal translations along the corresponding directions e_μ and e_ν.

The translation from pointwise commutativity (E5) to commutativity of infinitesimal translations relies on a local-triviality property: the algebraic structure assumed at x by (E5) extends coherently to a neighbourhood of x as a structure-preserving section of the local commitment-algebra bundle. This extension is implicit in (C3) locality — which assigns response at x dependent only on commitment structure in a neighbourhood of x, and therefore presupposes that the commitment algebra at x and the commitment algebras at neighbouring points x + dx are connected by a coherent local-triviality structure — but is not strictly contained in (E5) at a point. We flag this dependency explicitly: the derivation below uses (E5) at x together with the local-triviality structure implicit in (C3), and the conclusion is the commutativity of infinitesimal translations along independent directions in a neighbourhood of x.

Since the connection coefficients characterise the rate at which the local frame rotates under translation, the commutator of infinitesimal translations along e_μ and e_ν is precisely the torsion $T^\mu_{;\nu\lambda}$. Pointwise commutativity (extended to a neighbourhood by local triviality) therefore gives

$$T^\mu_{;\nu\lambda} = \Gamma^\mu_{;[\nu\lambda]} = 0 .$$

Why curvature is not similarly excluded. Curvature measures the holonomy of parallel transport around a *finite* region (equivalently, the second-order failure of an infinitesimal parallelogram to return the original frame), not the commutator of infinitesimal commitments at a single point. (E5) asserts pointwise commutativity at first order; it does not constrain second-order holonomy over a region. Curvature is therefore unconstrained by (E5) and is in fact required to encode gravitational response.

This separates torsion from curvature cleanly: torsion is excluded by the pointwise commitment-algebra structure (E5); curvature is unconstrained because it lives at a higher order where (E5) does not act.

Backup argument: (C6) minimality. Independently of (E5), torsion is excluded *at leading order* by (C6) in its operator-content-minimality formulation per §3: introducing a non-zero $T^\mu_{;\nu\lambda}$ carries independent operator content beyond the irreducible Lorentz-component structure required for universal coupling, since the universal coupling is supported by the Levi-Civita connection alone. (C6) therefore excludes torsion at leading order.

The (E5) argument and the (C6) argument do different work. (C6) gives leading-order torsion-freeness, sufficient for the leading-order theorems of §4 and §6. (E5) gives torsion-freeness at all orders, which is what Theorem 3's metric-uniqueness claim requires (the fundamental theorem of

Riemannian geometry produces a unique connection only for the strictly torsion-free case). The full strength of Theorem 3 therefore rests on (E5), not on (C6) alone.

9.2 From connection to metric

The committed structure at x is encoded in $\Phi_{\mu\nu}(x)$. In the locally homogeneous limit — where commitment density is approximately uniform on scales much larger than $\ell_\star$ but much smaller than the variation scale of $\Phi_{\mu\nu}$ — the transport map P must preserve $\Phi_{\mu\nu}$. We give the chain explicitly.

Status of the locally homogeneous limit. This limit is *not assumed at the fundamental level*. It arises as the coarse-grained regime in which the closure constraints (C1)–(C5) define transport consistently: at scales below $\ell_\star$ the substrate is not resolved (per (C1)), and at scales above the variation scale of $\Phi_{\mu\nu}$ the homogeneity assumption fails by definition. The intermediate window is the regime in which "parallel transport" is well-defined as a coarse-grained map between local commitment frames, and is therefore the natural regime in which to derive the metric structure. Outside this window, the metric description is not expected to hold — substrate-level corrections at sub-$\ell_\star$ scales, and inhomogeneity-driven corrections at super-$\Phi_{\mu\nu}$-variation scales, both lie outside the scope of Theorem 3 and are addressed in subsequent papers.

Chain (compatibility-or-contradiction). Suppose, for contradiction, that P does *not* preserve $\Phi_{\mu\nu}$ in the locally homogeneous limit. Then there exists a single committed structure C at event x such that:

1. *Local measurement.* An observer at the neighbouring event x′ measures C as a tensor with components $\Phi_{\mu\nu}(x')$ in their local frame, by the locally homogeneous assumption.
2. *Transported measurement.* An observer at x who transports C from x to x′ via P assigns components $P[\Phi_{\mu\nu}(x)]$ in the same local frame at x′.
3. *Non-preservation hypothesis.* By assumption, $P[\Phi_{\mu\nu}(x)] \neq \Phi_{\mu\nu}(x')$.
4. *Consequence.* The same physical commitment C is therefore assigned distinct componentwise values by two different observational routes (local measurement vs. parallel transport from a neighbour) at the same event x′.
5. *Contradiction with (C5).* Observer invariance requires that the components of a given physical commitment in a given local frame be observer-independent. Step 4 violates this directly.

The chain operates *post-emergence*: it presupposes that the local frame at x′ is well-defined, which is content of (E1) and the smooth-manifold structure. In the substrate-level pre-emergence picture, the local frame is itself derived. The chain is therefore a consistency argument given the emergent frame, not a substrate-level derivation; the latter is in Paper 3.

Therefore P must preserve $\Phi_{\mu\nu}$ in the locally homogeneous limit:

$\nabla\lambda\ \Phi\{\mu\nu\} = 0$ in the locally homogeneous limit.

Non-degeneracy of $\Phi_{\mu\nu}$. Before invoking the fundamental theorem, we must establish that $\Phi_{\mu\nu}$ is non-degenerate as a symmetric rank-2 form — i.e., that there is no non-zero vector v^μ with $\Phi_{\mu\nu} v^\nu = 0$. The argument is as follows.

Suppose for contradiction that $\Phi_{\mu\nu}$ were degenerate, with a non-zero v^μ in the null space. Then the commitment density along v vanishes: a commitment with directional structure aligned along v would produce zero contraction with $\Phi_{\mu\nu}$. By (C4) universality (per the definition in §3), the source $T_{\mu\nu}$ couples to *all* components of the commitment density.

The argument requires one bridge step worth flagging. (C4) as stated is a representation-theoretic claim — that R_F pairs non-trivially with all irreducible Lorentz components of R_T. That alone does not imply non-degeneracy: a degenerate $\Phi_{\mu\nu}$ would still couple non-trivially to T at the irrep level so long as some component of T outside Φ's null space is non-zero. The non-degeneracy claim genuinely requires an additional configuration-space assumption: that the matter sector includes commitments structured along *every* tangent direction, i.e., that $T_{\mu\nu}$ configurations of the form $v_\mu v_\nu$ for any $v \in T_x M$ are admissible. This is the content of (C4′) as labelled in §3 — a sub-input of (C4) supplying the configuration-space content beyond the bare irrep-level statement. Under (C4′), a direction along which $\Phi_{\mu\nu} v^\nu = 0$ would be invisible to the commitment-density coupling for the configuration $v_\mu v_\nu$, contradicting universality at the configuration level.

Therefore $\Phi_{\mu\nu}$ is non-degenerate, given (C4) at the irrep level and (C4′) at the configuration-space level.

Identification of $\Phi_{\mu\nu}$ as a metric. A non-degenerate symmetric rank-2 covariant tensor field is, by definition, a metric (with signature inherited from the corpus-level signature input flagged in §3). Therefore $\Phi_{\mu\nu}$ *is* a metric, and the torsion-free connection Γ compatible with $\Phi_{\mu\nu}$ (per §9.1) is, by the fundamental theorem of Riemannian geometry, the Levi-Civita connection of $\Phi_{\mu\nu}$.

The fundamental theorem is invoked here in its standard direction: *$\Phi_{\mu\nu}$ is a metric by definition* (non-degenerate symmetric rank-2 tensor), and *Γ is its Levi-Civita connection by the fundamental theorem*. The fundamental theorem does the work it standardly does — it determines the unique torsion-free metric-compatible connection given the metric.

What is non-definitional in this identification. A picky reader might worry that "metric emergence" reduces to a definitional unfolding: any non-degenerate symmetric rank-2 tensor *is* a metric, so calling $\Phi_{\mu\nu}$ a metric is not a derivation. The clarification: the *non-trivial* result of Theorem 3 is not the elementary set-theoretic fact that a non-degenerate symmetric rank-2 tensor satisfies the definition of a metric. It is that **$\Phi_{\mu\nu}$ is forced to be such a tensor by closure constraints** — symmetric by Theorem 1 (rank-selection from §4), non-degenerate by the (C4)-universality argument of §9.2, and admitting a torsion-free compatible connection by §9.1 plus (E5). These three properties are what the construction derives; the metric identification then follows by definition from the conjunction. The derivational content lives in compelling $\Phi_{\mu\nu}$ to satisfy the metric-defining properties, not in identifying the metric afterwards.

Normalisation. To match the standard convention of a dimensionless metric reducing to $\eta_{\mu\nu}$ in the vacuum (zero-commitment) limit, we define

$$g_{\mu\nu}(x) \equiv \Phi_{\mu\nu}(x) / \lambda_\star(x) ,$$

where $\lambda_\star(x)$ is the local commitment-density scale fixed by (C1). This is the unique normalisation for which $g_{\mu\nu}$ is dimensionless and reduces to $\eta_{\mu\nu}$ in vacuum.

Signature reminder. The signature of $g_{\mu\nu}$ is inherited from the Lorentzian tangent-space structure of M (corpus input (E2), per §3 and §12.2). The construction here does not derive signature, only the metric structure given the signature.

Theorem 3 (Metric Emergence). *Under (C1)–(C5) and (E5), with closure-consistent parallel transport of committed structure on the substrate:*

(i) ***Existence of connection.*** *There exists a torsion-free linear connection* $\Gamma^\mu{}_{\nu\lambda}$ *on the substrate.*

(ii) ***Existence of metric.*** *Compatibility of Γ with* $\Phi_{\mu\nu}$ *in the locally homogeneous limit determines, by the fundamental theorem of Riemannian geometry, a symmetric metric* $g_{\mu\nu}$ *(up to overall normalisation fixed by* $\ell_\star$*) such that Γ is the Levi-Civita connection of g.*

(iii) ***Uniqueness.*** *Both Γ and* $g_{\mu\nu}$ *are unique given the closure-consistent transport structure and the normalisation in (ii).*

The metric is therefore not a primitive ingredient of the theory and not a definition. It is the unique structure compatible with closure-consistent transport of committed records under (C1)–(C5) and (E5).

9.3 Lovelock applied to the geometric sector

Once $g_{\mu\nu}$ has been derived, its dynamics are constrained by Lovelock's theorem. In four dimensions, the most general second-order divergence-free symmetric tensor built from $g_{\mu\nu}$ and its derivatives is

$$\Lambda g_{\mu\nu} + \alpha (R_{\mu\nu} - (1/2) R g_{\mu\nu}) ,$$

i.e., the Einstein tensor plus a cosmological term. Coupling this to the stress-energy tensor $T^{(\Phi)}_{\mu\nu}$ of the Φ field (defined in §8.5) yields

$$G_{\mu\nu} + \Lambda g_{\mu\nu} = \kappa T^{(\Phi)}_{\mu\nu} ,$$

with κ fixed by the dimensional bookkeeping of §10.

Linearised vs full Einstein equation. Given the §8.1 background/perturbation split $\bar{\Phi}_{\mu\nu} = \lambda_\star \tilde{g}_{\mu\nu} + h_{\mu\nu}$, the stress-energy tensor $T^{(\Phi)}_{\mu\nu}$ is the *perturbation's* stress-energy on background $\tilde{g}$. The equation above is therefore the *linearised* Einstein equation governing perturbations of the background metric, not the full nonlinear Einstein equation governing the fundamental field $\bar{\Phi}$ directly. Recovery of the full nonlinear equation requires resumming the background/perturbation split — promoting $\tilde{g} + h$ to the full metric and reorganising the kinetic, source, and Einstein terms into their fully nonlinear forms. This is the standard linearised-gravity-to-full-GR resummation; in the present construction it is implicit and not carried out explicitly here. The linearised result is sufficient for the leading-order claims of this paper; the full nonlinear closure is treated in Paper 2.

The standard Lovelock hypotheses — covariance, second-order, divergence-free, symmetric rank-2 — are all derivable from (C1)–(C6) within the present construction: covariance from (C5), second-order from the leading-derivative-order truncation of (C3), divergence-free from the contracted Bianchi identity $\nabla^\mu G_{\mu\nu} = 0$ of the emergent geometry (a property of the LHS, distinct from the matter-side conservation identity $\nabla_\nu \mathcal{C}^{\mu\nu} = \mathcal{S}^\mu$ of §8.4), and symmetric rank-2 from Theorem 1. The Einstein form is therefore not invoked but follows from a standard theorem applied at the correct location in the dependency chain.

---

10. Dimensional Consistency and the Corpus Identity

The dimensional bookkeeping closes consistently when $\mathcal{C}^{\mu\nu}$ (the linear projection of $\Phi_{\mu\nu}$, properly the analogue of the trace-reversed metric perturbation) is distinguished from the stress-energy tensor $T^{(\Phi)}_{\mu\nu}$ (quadratic in Φ). With this distinction in place — already established in §8.5 — we now build the dimensional structure from canonical normalisation.

10.1 Canonical normalisation of $\Phi_{\mu\nu}$

The kinetic term in the master action (§8.2) is

$$\mathcal{L}_{\text{kin}} = -(1/2)\, \nabla_\lambda \Phi^{\mu\nu}\, \nabla^\lambda \Phi_{\mu\nu} \,.$$

The action $\int d^4x \sqrt{-g}\, \mathcal{L}$ is dimensionless in natural units ($\hbar = c = 1$), so the Lagrangian density satisfies $[\mathcal{L}] = M^4$. With $[\nabla] = M^1$, the kinetic term then requires $[(\nabla\Phi)^2] = M^4$, hence

$$[\Phi_{\mu\nu}] = M^1 \,.$$

This is canonical normalisation for a propagating tensor field on a 4D background.

10.2 Derived dimensions

From $[\Phi_{\mu\nu}] = M^1$ all other dimensional assignments follow:

- **Trace.** $[\Phi] = [g^{\mu\nu} \Phi_{\mu\nu}] = M^1$ (since $g^{\mu\nu}$ is dimensionless).
- **Linear projection coefficients.** $[a] = [b] = M^0$ (dimensionless), giving $[\mathcal{C}^{\mu\nu}] = M^1$.
- **Conservation source.** $[\nabla_\nu \mathcal{C}^{\mu\nu}] = M^1 \cdot M^1 = M^2$, hence $[\mathcal{S}^\mu] = M^2$.
- **Stress-energy tensor.** $[T^{(\Phi)}_{\mu\nu}] = M^4$ (standard energy density, quadratic in Φ).
- **Matter source.** $[T_{\mu\nu}^{matter}] = M^4$ (standard).
- **Mass parameter.** $[m] = M^1$ (in the $-(m^2/2)\, \Phi^2$ term of §8.2).

10.3 The coupling scale M_★

The finite-distinguishability scale $\ell_\star$ from (C1) supplies a UV scale

$M_\star \equiv \ell_\star^{-1}$, $[M_\star] = M^1$.

The leading matter coupling is then

$\mathcal{L}_{int} \supset (1/M_\star)\, \Phi^{\mu\nu}\, T_{\mu\nu}^{matter}$,

with dimensions $[(1/M_\star)\, \Phi\, T] = M^{-1} \cdot M^1 \cdot M^4 = M^4$ ✓. The factor $1/M_\star$ is the dimensional bridge between the canonically-normalised commitment field (dim M^1) and the matter stress-energy (dim M^4), exactly analogous to how $1/M_{Pl}$ bridges the canonically-normalised graviton and matter in standard gravity.

10.4 The corpus identity TPB[Φ] ~ |Φ|²

The conservation source $\mathcal{S}^\mu$ has dimension M^2. The corpus identity $TPB[\Phi] \sim |\Phi|^2$ requires that the source be quadratic in Φ at the schematic level (i.e., quadratic in field amplitude). With $[\Phi^2] = M^2$, dimensional matching at the scalar level is admissible.

The vector index requires explicit construction. $\mathcal{S}^\mu$ is a vector and $|\Phi|^2$ is a scalar — dimensional matching is necessary but not sufficient. We need a vector quadratic in $\Phi_{\alpha\beta}$ at dimension M^2. The available zero-derivative bilinear vector constructions from a symmetric $\Phi_{\alpha\beta}$ and metric $g_{\alpha\beta}$ alone are vacuous: every contraction of two symmetric rank-2 objects with the metric produces either too few or too many free indices to leave a single μ. The vector source $\mathcal{S}^\mu$ at the required dimension therefore requires a derivative on Φ to supply the index. With one derivative, the index count for a Φ^2 structure with one ∇ is $2\cdot2 + 1 = 5$ indices, all of which must contract pairwise plus one left free — admissible parity. The schematic structure is then $\nabla^\mu(\Phi\Phi)$ or similar, of dimension $M^1 \cdot M^1 \cdot M^1 = M^3$, which requires an explicit $(1/M_\star)$ factor to land at the M^2 source dimension.

The single bridge at the interaction vertex. Following this through: the relevant interaction term in $\mathcal{L}_{int}$ that produces the $|\Phi|^2$ conservation source by variation must itself respect dimensional consistency. A genuinely cubic Φ^3 operator with one derivative is impossible by index counting ($3\cdot2 + 1 = 7$ indices cannot pair). The minimal closed operator producing a dim-M^2 Φ-bilinear EOM contribution by variation is *quartic* in Φ with one derivative:

$\mathcal{L}_{int} \supset (1/M_\star)\, \Phi^{\mu\nu}\, \nabla_\mu \Phi_{\nu\alpha}\, \Phi^{\alpha\beta}\, \Phi_\beta{}^{;\nu}$

(or equivalent index-closed quartic structure). Dimensions: $(1/M_\star) \cdot M^1 \cdot M^2 \cdot M^1 \cdot M^1 = M^4$ ✓, with the $(1/M_\star)$ factor required for the vertex itself to be a valid Lagrangian operator. Variation with respect to $\Phi_{\mu\nu}$ produces an EOM contribution of schematic form $(1/M_\star) \cdot \Phi^2 \cdot \nabla\Phi$-type structure of dimension M^2, matching $\mathcal{S}^\mu$ directly without a second bridge factor.

The $(1/M_\star)$ that appears here is the *same* coupling scale as in §10.3's matter coupling $(1/M_\star)\, \Phi^{\mu\nu}\, T_{\mu\nu}^{matter}$. Both originate from the finite-distinguishability scale (C1) supplying the unique UV scale in the construction. There is no separate "self-interaction bridge" distinct from the "matter-coupling bridge" — the same $M_\star$ acts at both vertices, in line with minimality (C6): a different $M_\star$ for the two would have introduced an additional physical scale not present in (C1).

Honest realisation of the corpus identity. $TPB[\Phi] \sim |\Phi|^2$ is realised at the schematic level (quadratic in field amplitude) once the $(1/M_\star)$ coupling structure required by canonical normalisation is admitted. The realisation is *not* automatic in the sense of "no additional scale required": it requires the same $M_\star$ that mediates matter coupling, present at the interaction vertex itself rather than as a downstream bridge. The corpus identity survives in spirit, with its dimensional realisation tied to the same finite-distinguishability scale as everywhere else in the construction. This is a feature, not a bug: a single $M_\star$ appearing throughout is exactly what minimality (C6) requires.

The full coefficient algebra — including the dimensionless Wilson coefficient on the $(1/M_\star)$ vertex factor and its relation to the matter-coupling Wilson coefficient — is collected in Paper 2.

10.5 Summary

The dimensional structure that closes consistently:

Object	Dimension	Origin
$\Phi_{\mu\nu}$	M^1	Canonical normalisation of kinetic term
$\Phi \equiv g^{\mu\nu} \Phi_{\mu\nu}$	M^1	Trace
$\mathcal{C}^{\mu\nu}$	M^1	Linear projection of $\Phi_{\mu\nu}$
$\nabla_\nu \mathcal{C}^{\mu\nu} = \mathcal{S}^\mu$	M^2	Identity
$T^{(\Phi)}_{\mu\nu}$	M^4	Quadratic stress-energy from action variation
$T_{\mu\nu}^{matter}$	M^4	Standard energy density
$M_\star$	M^1	Finite-distinguishability scale from (C1)
Λ (cosmological term)	M^2	Standard
κ (Einstein coupling)	M^{-2}	Standard $1/M_{Pl}^2$

10.6 Resolution of the corpus dimensional issue

The $\mathrm{TPB}[\Phi] \sim |\Phi|^2$ identity, flagged at the corpus level as requiring resolution in the master-action sector, is realised at the schematic level (quadratic in field amplitude) within the present construction. With $[\Phi_{\mu\nu}] = M^1$ from canonical normalisation, and with $\mathcal{C}^{\mu\nu}$ (linear projection, dim M^1) distinguished from $T^{(\Phi)}_{\mu\nu}$ *(stress-energy, dim M^4), the dimensional bookkeeping closes. The realisation requires a single* $(1/M\bigstar)$ factor at the interaction vertex itself — the same scale that mediates matter coupling per §10.3 — bringing the quartic Φ-vertex operator to the required Lagrangian dimension M^4. Variation then produces a Φ-bilinear EOM contribution of dimension M^2 that sources the conservation identity directly, without an additional bridge. This is consistent with minimality: a single finite-distinguishability scale (C1) acts at all interaction vertices in the construction. The relationship between the dimensionless Wilson coefficient on this vertex and its analogue on the matter-coupling vertex is collected in Paper 2.

---

11. Predictions, Parametric Deviations, and Falsification

We extract three classes of quantitative deviation from General Relativity, then state a sharp falsification criterion.

11.1 Subleading nonlinear corrections

Nonlinear operators suppressed in §6 contribute corrections of order

$$\delta\mathcal{C}^{\mu\nu} / \mathcal{C}^{\mu\nu} \sim (\Phi / \Phi_\bigstar) \sim (\rho / \rho_\bigstar) ,$$

where $\rho_\bigstar \sim M_\bigstar^4$. For $M_\bigstar$ at the Planck scale, current astrophysical regimes have $\rho/\rho_\bigstar < 10^{-90}$. For $M_\bigstar$ at a lower scale — as suggested by the Two-Planck Principle — the suppression weakens and may become testable in compact-object mergers at strain amplitudes $h \sim 10^{-22}$ to 10^{-24}.

11.2 Memory effects from finite-distinguishability propagation

Higher-derivative corrections from (C1) introduce a non-Markovian kernel,

$$K(x, x') = K_{GR}(x, x') + (1/M_\bigstar^2)\, \Box\, K_{GR}(x, x') + \dots ,$$

manifesting as frequency-dependent dispersion of magnitude $\delta v/c \sim (\omega/M_\bigstar)^2$. At LIGO-band frequencies ($\omega \sim 10^2$ Hz) this is unobservably small for $M_\bigstar$ at the Planck scale but enters LIGO sensitivity for $M_\bigstar \lesssim 10^{-15}$ eV.

11.3 Trace contribution to the cosmological constant

The b $g^{\mu\nu}$ Φ term in $\mathcal{C}^{\mu\nu}$ contributes to the cosmological constant sector. Combined with the Two-Planck Principle estimate, the present construction predicts a trace contribution of fractional size

$$\delta\Lambda / \Lambda \sim b / (a + 4b) ,$$

providing a corpus-internal consistency check.

11.4 Falsification criteria

The construction admits the following sharp falsifiers, each of which, if observed, would refute one or more of its core claims:

(F1) Detection of antisymmetric rank-2 phenomenology mimicking gravity. Theorem 1 establishes on representation-theoretic grounds that an antisymmetric rank-2 field cannot couple to the symmetric stress-energy tensor $T_{\mu\nu}$ as a Lorentz scalar bilinear interaction. As a strict mathematical statement, this is unfalsifiable — the Schur intertwiner does not exist, and no observation can change that. The genuinely empirical falsifier in this sector is the following: detection of a fundamental rank-2 field whose phenomenology mimics gravitational response — universal coupling to all matter sources with equivalence-principle-like behaviour — but which exhibits *antisymmetric* features in its propagation, polarisation, or coupling structure.

Concretely, the LVK general-polarisation analysis framework [1] — applied at population level in the GWTC-3 tests-of-GR analysis [2] — decomposes a detected strain into a basis of six possible polarisation modes: two transverse-traceless tensor modes (+, ×), two vector modes (x, y), and two scalar modes (breathing and longitudinal). Standard GR predicts only the two tensor modes; alternative gravity theories predict one or more of the additional four. The GWTC-3 analysis uses Bayesian model comparison (purely tensor vs. purely vector vs. purely scalar vs. mixed) and reports no preference for non-tensor modes across the catalogue.

The current observational situation has two structural limitations that constrain the strength of any falsifier in this sector. First, the Eardley et al. result [3] establishes that at least five non-coaligned differential-arm detectors are required to break the degeneracies among all five non-degenerate polarisation modes; the existing LVK network does not yet meet this threshold for transient sources, so present constraints rely on either Bayesian model comparison among restricted hypotheses or on tests using continuous sources or multi-messenger events with electromagnetic counterparts (such as GW170817-class binary neutron star mergers with localisation from associated GRBs). Second, an antisymmetric rank-2 source coupling — distinct from the standard scalar-tensor and vector-tensor alternatives that drive most current analyses — would manifest as a specific population-level statistical signature, particularly in the sky-position dependence of the strain transfer function under cross-event correlation. The morphology-independent analyses of [4] provide the technical infrastructure to test for such generic non-tensorial content.

A clean detection of antisymmetric-rank-2 phenomenology — distinguishable from scalar-tensor and vector-tensor alternatives by its specific sky-position-dependent signature in cross-event polarisation correlations — would falsify the source-identification step of (C4), even though it would not falsify Theorem 1 mathematically. Quantitative bounds tighten by roughly an order of magnitude with the addition of LIGO-India to the network and again with next-generation observatories (Einstein Telescope, Cosmic Explorer, LISA), at which point the five-detector Eardley threshold is fully met and the falsification potential becomes correspondingly sharper.

(F2) Absence of $(\omega/M_\star)^2$ dispersion below the suppression scale. Equation §11.2 predicts a definite parametric dispersion of gravitational waves. The construction is testable concretely as follows: in the LIGO–Virgo–KAGRA (LVK) network, a binary neutron star merger of GW170817 class — observed in the gravitational-wave band 30–300 Hz with a multimessenger electromagnetic counterpart over a propagation distance of order 40 Mpc — provides a cross-frequency timing constraint on $\delta v/c$ at the observed band. The framework predicts a frequency-dependent arrival-time dispersion of order $\delta t(\omega) \sim D\cdot(\omega/M_\star)^2/c$ across the LVK band; the absence of such dispersion in current and future BNS events of this class, at sensitivities tighter than the parametric prediction, would force $M_\star$ above the assumed scale or, when combined with the corpus constraints from the Two-Planck Principle, falsify the construction. The next-generation observatories (Einstein Telescope, Cosmic Explorer) will tighten this bound by roughly an order of magnitude in $M_\star$ via improved low-frequency sensitivity.

(F3) Observation of inverse-square-law deviations at densities $\rho \ll \rho_\star$ that do not follow the parametric form §11.1. The construction predicts that nonlinear corrections to GR scale as $(\rho/\rho_\star)$. A measured deviation from inverse-square scaling at low density not fitting this parametric form — equivalently, deviations whose density-scaling exponent is incompatible with the leading-order linear ansatz — would falsify either Theorem 2 or the suppression-scale argument of §6.

(F4) Non-vanishing antisymmetric component of the gravitational connection. Theorem 3 derives the metric as a symmetric rank-2 field via the fundamental theorem of Riemannian geometry from a torsion-free connection, with torsion-freeness resting on the supplementary input (E5) (pointwise commitment commutativity). Pointwise commutativity is itself an algebraic claim about the substrate-level commitment algebra; as a strictly algebraic statement it is unfalsifiable via direct observation of the substrate, since the substrate is below the resolution scale $\ell_\star$. The empirical falsifier in this sector is therefore (E5)'s observational consequence: detection of a fundamental torsional component of the gravitational connection — measurable in principle through tests of the equivalence principle for spinning bodies (Hayasaka–Takeuchi-type protocols at improved sensitivity, or equivalent cold-atom interferometry) — would falsify the (E5)-based derivation of torsion-freeness in §9.1, equivalently forcing a non-trivial pointwise commitment-algebra structure in the substrate (the substrate algebra would be non-abelian at first order in the resolution scale $\ell_\star$). This is a sharper empirical handle than "torsion detection falsifies §9.1": it specifies which input fails (the algebraic claim of (E5)) and in which direction it must be revised (toward a non-abelian substrate algebra at first order). The framing parallels (F1): the underlying input is algebraic, the falsifier is the observational consequence.

The most practically accessible of these is (F2). The most structurally consequential is (F1) — not because it would falsify Theorem 1 mathematically (Schur's lemma is unfalsifiable), but because it would falsify the source-identification step in (C4) and force a reformulation of the universality assumption.

---

12. Logical Dependency Structure

We give the explicit logical structure of the construction, separating axioms from supplementary inputs from derived results from conditional claims. The arrows ⇒ denote logical entailment; bracketed labels mark which constraints are invoked.

12.1 Primary axioms (closure constraints)

- **(C1)** Finite distinguishability — primitive
- **(C2)** Irreversible commitment — primitive
- **(C3)** Locality — primitive
- **(C4)** Additivity, including universality of the source — primitive
- **(C5)** Observer invariance — primitive
- **(C6)** Minimality, formalised — primitive

These are the only inputs derived from nothing else within the closure framework.

12.2 Supplementary inputs (not derived from C1–C6)

The construction is not a pure derivation from the closure axioms alone. Five supplementary corpus inputs are required. They are not arbitrary additions: each is the continuum-limit expression of a structural feature already implicit in record-based physics, and each is motivated below as the minimal bridge assumption required to pass from the substrate-level record ontology to a smooth post-emergence gravitational theory.

(E1) Smooth 4-manifold M. The smooth manifold M is not postulated as fundamental. It is the continuum-limit representation of the substrate commitment network obtained when discrete commitment events are coarse-grained over scales much larger than $\ell_\star$ — directly analogous to the way smooth fluid fields (density, velocity, pressure) emerge from molecular water once the coarse-graining scale contains many molecules. Its four-dimensional character can be motivated by separating one irreversible ordering direction, supplied by the substrate sequence of commitments, from three independent relational degrees of freedom required for stable localisation, propagation, and isotropic dilution of influence.

The strongest available argument for *three* spatial dimensions specifically — beyond informal stability and propagation considerations — ties dimensionality directly to the result the paper is constructing. In a manifold with d spatial dimensions, the divergence theorem for a conserved current of fixed source strength gives field strength scaling as $r^{-(d-1)}$. Only at $d = 3$ does this yield the inverse-square law r^{-2}, which is the precise scaling required for stable long-range coupling sourced by a conserved energy-momentum tensor — equivalently, the unique scaling at which orbital and bound-state structure are stable rather than collapsing or escaping. Conservation (from (C5) via Noether) plus locality (C3) plus the requirement of stable long-range gravitational structure therefore selects d = 3 spatial dimensions as the unique value compatible with the construction's downstream content. Fewer than three (d = 2 gives r^{−1} field strength, with logarithmic potential and no stable bound orbits; d = 1 gives constant field strength, with no spatial fall-off at all) fails stable long-range structure; more than three (d = 4 gives r^{−3}, with no stable bound orbits per Bertrand's theorem and the well-known higher-dimensional instability of Kepler problems) fails the same requirement.

This sharpens the dimensionality argument from "plausible minimality" to "structural necessity given the conservation/locality structure (C3)+(C5) and the requirement of stable long-range coupling that the construction's downstream gravitational content presupposes." Note that this argument is not fully internal to the paper: it presupposes that the construction is to produce stable long-range gravitational coupling, which is content of the construction's purpose rather than of (C1)–(C6) directly. (E1) thus remains a supplementary input — but its dimensionality clause is now tied to the construction's downstream physics rather than to informal minimality.

Thus 3+1 is the minimal smooth continuum compatible with local propagation, stable relational structure under inverse-square dilution, and a single irreversible ordering parameter; it is taken as input here, not derived. M is best understood as the minimal effective arena in which the record network admits smooth local fields, not as a primitive container. (E1) is presupposed by (C3); rigorous derivation is in the time-emergence sector of the corpus.

(E2) Lorentzian (3,1) signature. The signature follows naturally once the emergent continuum distinguishes irreversible commitment order from reversible relational extension. The commitment sequence supplies one direction with causal ordering and non-reversible structure; spatial directions encode reversible comparison among coexisting records. A Euclidean signature would erase this asymmetry — treating all four directions as equivalent and conflicting with irreversible commitment. Multiple time-like directions would create independent commitment orderings, risking causal ambiguity and undermining well-posed initial-value evolution. The signature (3,1) is therefore the minimal continuum signature compatible with one irreversible record-ordering parameter and three relational degrees of freedom, *given* an additional structural property of the substrate-to-continuum coarse-graining map analysed in Appendix A.2 and named (E2b) below. (E2) is built into (C5) via the SO(3,1) covariance reduction; the whole §4 representation-theoretic argument is conducted in irreps of SO(3,1).

(E2b) Topological invariance of the ordering distinction. A.2 establishes that the reduction of (E2) to (E1) plus the substrate-level ordering contained in (C2) requires a separate structural property: the coarse-graining map π must encode the ordering distinction in the tangent-space inner product *topologically* (via sign) rather than via magnitude (via anisotropic Riemannian rescaling). This is not a consequence of (E1), (C2), or any other input above — it is a separate substrate-level property of π. We name it explicitly:

Under coarse-graining $\pi: \mathcal{N} \to M$, the substrate-level ordering distinction at each event $x \in M$ is encoded in the tangent-space inner product as a topologically invariant feature (sign-based) rather than as a magnitude-based feature.

The logical structure is then: **(E2) follows from (E1) + (C2)+(E4) + (E2b)**. (E2) is not an independent input given (E2b), but (E2b) is itself a substrate-level property whose verification requires the explicit form of π and is therefore deferred to Paper 3. Naming (E2b) explicitly clarifies that the *real* substrate-level assumption underlying signature is the topological-invariance property, not signature itself. The dynamical motivation for (E2b) — that only indefinite (and therefore sign-encoded) signature supports hyperbolic field equations consistent with irreversibility per (C2) — is given in A.2.

In the supplementary-inputs accounting, (E2b) replaces (E2) as the genuinely independent input. We retain (E2) in the §12.2 list as the most familiar formulation for readers approaching the construction from a standard-physics direction; the (E2b) name is used in Appendix A and in the dependency map above.

(E3) EFT power-counting. EFT power-counting is the continuum expression of finite distinguishability, not an independent modelling assumption. Since (C1) imposes a finite resolution scale $\ell_\star$ — i.e., no observable may depend on sub-$\ell_\star$ structure — the continuum theory cannot contain unsuppressed sensitivity to arbitrarily short-distance structure. Operators involving additional derivatives or higher powers of $\Phi_{\mu\nu}$ must therefore appear with powers of the cutoff scale $M_\star = \ell_\star^{-1}$ unless a new low-energy degree of freedom appears. The EFT hierarchy is in this sense the natural low-energy bookkeeping implied by finite distinguishability. The suppression-scale argument of §6 and §11 presupposes (E3). Making the connection between (C1) and (E3) fully rigorous — i.e., deriving NDA power-counting from substrate-level coarse-graining rather than motivating it from (C1) — is the subject of Paper 3, where the matching procedure produces (E3) as a derived rather than assumed input.

(E4) Substrate-level ordering of commitment events. A substrate-level ordering of commitments is required by the notion of irreversible record formation itself. A commitment event is not merely an element of a set; it is an event whose occurrence constrains subsequent admissible extensions of the record. Without an ordering relation, there is no distinction between prior record and candidate future extension, and hence no well-defined irreversible commitment process. Crucially, this ordering need not be continuous time — it is a *pre-temporal partial order* of fact formation, weaker than time, from which the emergent temporal coordinate is recovered only in the coarse-grained manifold limit. The notion of "history" in (C2) — along which $\mathcal{S}[\Phi]$ is non-decreasing — refers to this substrate-level ordering, not to a pre-existing temporal coordinate. (E4) is presupposed by (C2); derivation in the time-emergence sector of the corpus.

(E5) Pointwise commitment commutativity. This input must be read narrowly. It does *not* claim that all physical operations commute, nor does it deny quantum noncommutativity. It states only that infinitesimal contributions to the same local commitment density compose commutatively at first order: if two infinitesimal commitments dC_1 and dC_2 occur at the same substrate point, then the accumulated local density is independent of whether dC_1 is recorded before dC_2 or vice versa,

$$dC_1 \circ dC_2 = dC_2 \circ dC_1 \,.$$

If this failed, the substrate would retain an additional order-sensitive internal structure at a point, producing torsion-like degrees of freedom or internal-gauge-like residue. Such structure would introduce additional degrees of freedom not required for the universal gravitational coupling identified in §4 and would therefore violate (C6) minimality unless empirically required. (E5) is thus the *minimal gravitational-sector* assumption: local record density is abelian at first order, while any noncommutative residue belongs to subleading torsional or gauge-like sectors and is not part of the universal gravitational response. (E5) is required for the torsion-freeness derivation of §9.1 and hence for the full strength of Theorem 3 (metric uniqueness); rigorous derivation is deferred to Paper 3.

We flag this is the construction's largest substrate-level structural assumption: (E5) is the input that excludes torsion at the substrate level, and a sharp reviewer is right to note that torsion-killing happens here rather than as a derived consequence of more primitive axioms. The (C6)-based exclusion of non-commutative substrate structure (additional degrees of freedom not required for universal gravitational coupling) gives a structural reason to take (E5) as the minimal choice, but does not derive it from elsewhere — that derivation requires the substrate-level commitment algebra construction of Paper 3.

Coherent bridging. The five supplementary inputs are not an unrelated list of assumptions. Each is the continuum-limit expression of a structural feature already implicit in record-based physics: a smooth four-dimensional manifold is the coarse-grained limit of many discrete commitments; Lorentzian signature reflects the distinction between irreversible ordering and reversible relational extension; EFT power-counting expresses finite distinguishability as a UV cutoff; substrate ordering is required for irreversible record formation; and pointwise commutativity states that local record density is additive at first order. They form a coherent bridge from the discrete substrate to the post-emergence tensorial theory. They are not derived in this paper, but they are not free either: they are the minimal bridge assumptions required to pass from the substrate-level record ontology to a smooth post-emergence gravitational theory. Appendix A takes a step closer to derivation for three of the five and shows that the *substantive* supplementary content effectively reduces to approximately three core inputs — the five-input presentation here is retained for readability and for the dependency-tracking work it does in §12.

Dependency map. The primary dependencies among constraints, supplementary inputs, theorems, and downstream sections are:

Dependent	Depends on
(C2)	(E4)
(C3)	(E1)
(C5)	(E1), (E2)
§6 / §11 suppression-scale argument	(E3)
§9.1 torsion-freeness	(E5)
Theorem 1	(C1)–(C6), (E2)
Theorem 2	(C3)–(C6), (E3)

Dependent	Depends on
Theorem 3	(C1)–(C5), (C4′), (E1), (E5); independent of (C6) and Theorems 1, 2
§8 master action	Theorem 3 (presupposes emergent metric)
§9.3 Einstein structure	Theorem 3, Lovelock's theorem

Without (E5), only (C6)-driven leading-order torsion-freeness is available, which is sufficient for §4 and §6 but insufficient for the full uniqueness in Theorem 3. The independence of Theorem 3 from (C6) and Theorems 1–2 is highlighted in §12.5; it means the metric-emergence argument survives even if the rank-selection or linear-uniqueness machinery is weakened.

12.3 Derived results

- **Lemma (Schur-type, §4.2)**, derived from (C3), (C4), (C5). Scope: bilinear in F and T at zero-derivative order in F.
- **Corollary (containment-to-equality, §4.2)**, derived from the Lemma plus (C6).
- **Theorem 1 (Rank Selection, §4.4)**, derived from the Lemma + Corollary + Lorentzian signature input. Depends on all of (C1)–(C6) plus the signature input.
- **Symmetry of $\mathcal{C}^{\mu\nu}$ (§7)**, derived from Theorem 1.
- **Theorem 3 (Metric Emergence, §9)**, derived from (C1)–(C5) plus the supplementary input (E5) (pointwise commitment commutativity, required for torsion-freeness) plus the fundamental theorem of Riemannian geometry. Includes the non-degeneracy argument from (C4)+(C4′) — universality at the irrep level (C4) plus configuration-space inclusion (C4′). **Does not require (C6) and does not require Theorems 1 or 2.** Without (E5), only leading-order torsion-freeness from (C6) is available, which is sufficient for §4 and §6 but insufficient for the full uniqueness statement of Theorem 3.
- **Master action (§8)**, the post-emergence effective description. **Presupposes Theorem 3** (uses $g_{\mu\nu}$ as background; uses $\sqrt{-g}$ measure and $g^{\mu\nu}$ contractions). Constructed to be consistent with Theorem 2's linear projection structure for $\mathcal{C}^{\mu\nu}$.
- **Theorem 2 (Linear Uniqueness, §6)**, derived from (C3), (C4), (C5), (C6) within the linear leading-derivative-order (zero-derivative-in-Φ) ansatz. Theorem 2 fixes the form of the linear projection $\mathcal{C}^{\mu\nu}$; it is independent of the master action's specific form but consistent with it. Requires the EFT power-counting supplementary input to bound subleading corrections.
- **Conservation identity (§8.4)**, derived from diffeomorphism invariance of $\mathcal{S}$ (Noether's second theorem). Requires the master action of §8, hence presupposes §9.
- **Distinction of $\mathcal{C}^{\mu\nu}$ from $T^{(\Phi)}_{\mu\nu}$ (§8.5)**, derived from the action structure. Required for dimensional consistency of the construction.
- **Dimensional bookkeeping (§10)**, derived from canonical normalisation of the kinetic term in §8.2, hence presupposes §8 and (transitively) §9.
- **Einstein structure (§9.3)**, derived from Theorem 3 plus Lovelock's theorem applied to the geometric sector. Requires the standard Lovelock hypotheses, all of which are themselves derivable from (C1)–(C6) within the present construction.

12.4 Conditional / not derived in this paper

- **Numerical values of a, b, κ_1, κ_2, m, $M_\star$.** These appear as undetermined coefficients in the master action. Their determination requires empirical input from the matter sector and is the subject of Paper 4.
- **Full nonlinear closure beyond leading order.** The linear ansatz of Theorem 2 is parametrically justified at low density but not derived for arbitrary Φ.
- **Specific form of $\mathcal{L}_{int}$.** Only its schematic form is fixed by dimensional bookkeeping; the full structure depends on the matter sector.
- **Quantitative form of subleading corrections.** §11 provides parametric estimates; full coefficients await Paper 2.
- **Coarse-graining map from microscopic substrate to $\Phi_{\mu\nu}$.** Theorem 3 treats this at the level of structure rather than computation; full derivation is in the master-action completion paper.
- **Pre-emergence formulation of the action.** The action of §8 is the post-emergence effective description; the pre-emergence formulation (without $g_{\mu\nu}$) is the subject of Paper 3.

12.5 Independence and ordering relations

Three structural relations are worth stating explicitly:

- **§9 is logically prior to §8.** The master action presupposes the metric derived in Theorem 3. The order of presentation (§8 before §9) is for expository clarity, not logical dependency.
- **Theorem 3 is independent of Theorems 1 and 2 (but not of (E5)).** Metric emergence requires only (C1)–(C5) plus the supplementary input (E5) and the local-triviality structure implicit in (C3); it does not require (C6), Theorem 1, or Theorem 2. This means that even if Theorems 1 or 2 were weakened or replaced, the metric-emergence argument would survive — provided (E5) holds.
- **Falsifiers (F1)–(F4) target distinct sectors of the construction.** (F1) targets the source-identification step in (C4); (F2) targets the suppression-scale argument of §6 and §11; (F3) targets Theorem 2; (F4) targets the (E5)-grounded torsion-free derivation in §9.1. The construction therefore has multiple independent failure modes, each empirically distinguishable.

12.6 Honest accounting

The above dependency structure exposes that the paper's central derivations rely on:

- six primary closure axioms (C1)–(C6),
- five supplementary corpus inputs (E1)–(E5): smooth 4-manifold, Lorentzian signature, EFT power-counting, substrate-level ordering, pointwise commitment commutativity,
- one geometric theorem from outside VERSF (the fundamental theorem of Riemannian geometry),
- one structural theorem from outside VERSF (Lovelock's theorem in 4D).

The construction is therefore not a derivation of gravity from closure axioms *alone*; it is a derivation of gravity *on the emergent manifold*, given the closure axioms plus a small set of supplementary corpus inputs that act as the minimal bridge from the substrate-level record ontology to the post-emergence smooth theory, plus two standard external theorems. The bridge inputs are not arbitrary: each is the continuum-limit expression of a structural feature already implicit in record-based physics, motivated in §12.2. A construction running from substrate directly to dynamics in a single paper is not attempted here and would naturally form part of Paper 3. This is, we believe, the most that can be honestly claimed for any derivation of geometric gravity from foundational principles, and the explicit accounting above is intended to be a model for the rest of the corpus.

Effective minimality. A reader concerned with minimality of the assumption set should note Appendix A.4: the substantive content of the five supplementary inputs effectively reduces to approximately three core inputs — (E1) the smooth 4-manifold, (E3b) Wilson-coefficient naturalness, and (E5) pointwise commitment commutativity — with (E2) reducible to (E1) modulo (E2b), the substrate-level topological-invariance condition on the coarse-graining map (named explicitly in §12.2 (E2) entry and grounded dynamically in A.2); (E3a) derivable from (C1)+(C6); and (E4) contained at the level of existence within (C2). The five-input presentation in the main body is retained for readability; the deeper analysis is in the appendix.

---

13. Scope and Limitations

Post-emergence scope. The construction is *post-emergence*: it derives the leading-order tensorial dynamics on the already-emerged 4-manifold and the metric structure on it, but does not derive the manifold's existence, dimension, smooth structure, or signature. These are corpus inputs (E1)–(E2) per §12.2. A construction running from substrate directly to dynamics in a single paper is not attempted here and would naturally form part of Paper 3 of the master-action roadmap.

The uniqueness theorems of §4 and §6 are *constrained* uniqueness theorems. Theorem 1 is uniqueness within the constraint hierarchy (C1)–(C6) plus the Lorentzian signature input (E2), with (C6) formalised as in §3, and within the class of finite-dimensional irreducible representations of SO(3,1). The Schur-type lemma underpinning Theorem 1 is scoped to bilinear interactions at zero-derivative order, with higher-multilinear and higher-derivative couplings excluded at leading order by (C6). Theorem 2 is uniqueness within the linear, local, leading-derivative-order (zero-derivative-in-Φ) ansatz, with nonlinear and higher-derivative corrections suppressed by EFT power-counting in ($\Phi/\Phi_\star$) and ($\partial/M_\star$) — itself a supplementary input (E3) as flagged in §12.2. Theorem 3 is uniqueness up to overall normalisation, with the normalisation fixed by $\ell_\star$, and depends on the non-degeneracy argument from (C4)+(C4′) — universality at the irrep level plus configuration-space inclusion.

The values of the coefficients a, b, κ_1, κ_2, m, $M_\star$ are not derived in this paper; they appear as undetermined coefficients in the master action and require empirical or matching input from the matter sector for their determination.

The Einstein-emergence argument of §9.3 invokes Lovelock's theorem applied to the *geometric* sector (LHS of the Einstein equation), not to the commitment current. The standard Lovelock hypotheses are recovered from VERSF closure as indicated in §9.3.

The antisymmetric rank-2 sector, ruled out as a candidate for *gravitational* response in §4.3, is not ruled out as a physical field. Its potential role in a richer programme (torsional, Kalb–Ramond, electromagnetic-like) is left open.

Signature. The Lorentzian (3,1) signature is presupposed via the SO(3,1) framing of (C5) and listed as supplementary input (E2). Its derivation from the irreversibility-emergent-time programme is the subject of separate work in the corpus.

Manifold structure. The smooth 4-manifold M on which the dynamics live is presupposed (input (E1)) and not derived in this paper. Its derivation from substrate dynamics is in the time-emergence sector of the corpus.

---

14. Conclusion

Six structural claims have been established under the constraint hierarchy (C1)–(C6), supplemented by five corpus inputs (smooth manifold M, Lorentzian signature, EFT power-counting, substrate-level commitment ordering, and pointwise commitment commutativity — (E1)–(E5)) which together serve as the minimal bridge from the substrate-level record ontology to the post-emergence smooth theory, and two standard external theorems (the fundamental theorem of Riemannian geometry, Lovelock's theorem in 4D).

Scalar closure is insufficient for gravity, since it cannot encode pairwise directional structure or couple universally to the trace-free sector of the committed energy-momentum tensor.

Symmetric rank-2 response is the unique minimal admissible irrep of SO(3,1) admitting universal coupling to the committed energy-momentum source, by the Schur-type lemma of §4.2 and Theorem 1. Vector and antisymmetric rank-2 candidates are ruled out at leading order by representation-theoretic impossibility (vector strictly so at zero-derivative order, with one-derivative couplings excluded by (C6)), not by sign or coupling-convention arguments.

The leading-order record current is uniquely $\mathcal{C}^{\mu\nu} = a\,\Phi^{\mu\nu} + b\,g^{\mu\nu}\,\Phi$ within the linear leading-derivative-order ansatz. The current is a *linear projection* of $\Phi_{\mu\nu}$ (analogue of the trace-reversed metric perturbation), distinct from the stress-energy tensor $T^{(\Phi)}_{\mu\nu}$ (quadratic in Φ) that sources the Einstein equation. This distinction is essential: the two objects play different roles in the construction.

The master action of §8 — in its post-emergence effective form, presupposing the metric of Theorem 3 — is consistent with the linear-projection structure of $\mathcal{C}^{\mu\nu}$, closes the conservation identity by Noether's second theorem, and admits a dimensional bookkeeping that follows from canonical normalisation. The corpus identity TPB[Φ] ~ $|\Phi|^2$ is realised at the schematic level via the same $M_\bigstar$ bridge factor that mediates matter coupling.

The metric $g_{\mu\nu}$ emerges as the unique structure compatible with closure-consistent transport of committed records: torsion-freeness from the supplementary input (E5) (pointwise commitment commutativity), non-degeneracy from (C4) universality, and identification as a metric by definition, with the Levi-Civita connection then identified by the fundamental theorem of Riemannian geometry. The metric is not a primitive postulate.

The Einstein field equation emerges as the unique low-derivative covariant closure of the *geometric* sector via Lovelock's theorem applied at the correct location in the dependency chain — to the LHS of the equation, not the RHS, sourced by $T^{(\Phi)}_{\mu\nu}$ not by $\mathcal{C}^{\mu\nu}$.

Gravity, in the VERSF programme, is therefore not an independent force. It is the leading-order tensorial closure of record dynamics under the constraints required for a universe of irreversible facts — with the rank, symmetry, metric, and dynamical form of the response sector each *selected as the unique admissible structure within the stated constraint class*, rather than postulated. The construction's reliance on supplementary inputs and standard external theorems is laid out explicitly in §12, in the spirit of being honest about what is genuinely derived and what is genuinely assumed.

References

[1] M. Isi and A. J. Weinstein, "Probing gravitational wave polarizations with signals from compact binary coalescences," arXiv:1710.03794 (2017).

[2] R. Abbott et al. (LIGO Scientific Collaboration, Virgo Collaboration, and KAGRA Collaboration), "Tests of General Relativity with GWTC-3," arXiv:2112.06861 (2021).

[3] D. M. Eardley, D. L. Lee, A. P. Lightman, R. V. Wagoner, and C. M. Will, "Gravitational-Wave Observations as a Tool for Testing Relativistic Gravity," Phys. Rev. Lett. 30, 884 (1973).

[4] M. Isi et al., "A morphology-independent test of the mixed polarization content of gravitational wave signals from compact binary coalescences," arXiv:2105.01521 (2021).

Appendix A: Toward Substrate-Level Reductions of the Supplementary Inputs

The five supplementary inputs (E1)–(E5) of §12.2 are flagged as corpus dependencies whose rigorous derivation lies outside this paper. The motivations given in §12.2 are plausibility arguments, not derivations. This appendix takes a step closer to derivation for three of the five, in increasing order of structural cost:

- **A.1.** A definitional reduction of (E4) to (C2): the substrate-level ordering of commitment events is *contained* in (C2)'s notion of irreversible accumulation, not added to it.
- **A.2.** A plausibility derivation of (E2) Lorentzian signature from (E1) plus the substrate ordering, formalising the §12.2 motivation as a structural argument.
- **A.3.** A partial derivation of (E3) EFT power-counting from (C1) finite distinguishability via standard decoupling, with explicit acknowledgement of the residual Wilson-coefficient gap.

We do not attempt analogous reductions for (E1) (the smooth 4-manifold itself, whose derivation properly belongs to the time-emergence and K=7-closure sectors of the corpus) or (E5) (pointwise commitment commutativity, where a derivation from (C1)–(C5) is not available within the present construction; (E5) is therefore taken as an independent supplementary input). Re-opening (E5) here would either rediscover a (C2)+(C3) framing whose limitations have been discussed in §9.1, or require importing substrate-algebra machinery beyond the paper's scope.

The three reductions below are not intended to eliminate the supplementary inputs but to clarify their structure: which inputs are definitional unfoldings of axioms already accepted, which are plausibility-derivable from substrate features, and which are residual gaps awaiting the matching procedure of Paper 3. The architecture of §12.2 is unchanged; the appendix sharpens the picture of *why* the inputs are there.

A.1 (E4) as a definitional unfolding of (C2)

Claim. (E4) — the existence of a substrate-level ordering of commitment events — is a definitional consequence of (C2) irreversibility, not an additional input. The two should be read as a single axiom rather than as constraint plus supplementary input.

Argument. (C2) states that the commitment-density functional $\mathcal{S}[\Phi]$ is non-decreasing along admissible histories. This statement is not well-formed unless "non-decreasing" is interpreted relative to some ordering on the events along which $\mathcal{S}$ is evaluated. The ordering is therefore presupposed by the very assertion of (C2): it is what "non-decreasing" means.

Concretely, suppose (C2) held with respect to *no* ordering relation on commitment events. Then the assertion "$\mathcal{S}[\Phi]$ is non-decreasing along histories" would have no content: there would be no notion of "earlier" or "later" along which to compare values of $\mathcal{S}$. The statement would reduce to

"S is single-valued on the set of commitment events," which is just well-definedness, not irreversibility.

Equivalently: irreversibility *is* the assertion that the inverse direction of an ordering is not admissible. To assert irreversibility is to assert that there is an ordering whose reverse is forbidden. Without an ordering, "irreversible" is contentless.

Consequence for the dependency structure. (E4) does not add anything to (C2) at the level of *existence* of the ordering; it makes explicit a structure that (C2) already presupposes. The supplementary-inputs list of §12.2 could in principle be tightened from five to four by treating the existence of the ordering as contained in (C2). With that reading, the count is:

- six closure axioms (C1)–(C6),
- four supplementary corpus inputs (E1), (E2), (E3), (E5),
- two external theorems.

A subtlety worth flagging. What (C2)+(E4) collectively assert is the existence of an ordering. They do *not* assert that the ordering is total, that it is locally finite, that it is dense, or that its coarse-graining yields a continuous temporal coordinate. Each of these stronger properties is required at various downstream points in the construction (the locally homogeneous limit of §9.2 implicitly needs density; the matching to a smooth temporal coordinate needs the coarse-graining limit to behave well). These stronger properties belong to the time-emergence sector of the corpus and are not extracted here. (C2)+(E4) supplies only the bare ordering relation; the structural properties of that ordering — the residual content of "(E4)" beyond what (C2) alone gives — are inherited from (E1) (the smooth manifold's temporal coordinate) and the corpus-level emergence story.

In summary: the *existence* of the substrate-level ordering is contained in (C2). The *coarse-graining of that ordering into a smooth temporal coordinate* is content of (E1). What was previously labelled (E4) splits cleanly into these two pieces, with no residual independent content.

Recommendation. Rather than dropping (E4) as a separate input, the cleanest reading is that *(E4) can be understood as contained within (C2) at the level of existence, while its structural properties — totality, density, coarse-graining to a continuous temporal coordinate — remain tied to (E1)*. Future revisions might add a single clarifying sentence to (C2) ("...along the substrate-level ordering whose existence is contained in this constraint") and let the smooth-temporal-coordinate aspect be carried by (E1). The dependency table simplifies; the supplementary-inputs list is conceptually shortened; nothing in the construction breaks. We retain (E4) as a separate entry in §12.2 for readability, with the understanding that the *substantive* content of (E4) beyond what (C2) and (E1) jointly supply is residual and small.

A.2 (E2) Lorentzian signature from (E1) plus substrate ordering

Claim. Given (E1) — the smooth 4-manifold M as the coarse-graining limit of the substrate commitment network — and the substrate-level ordering contained in (C2) per A.1, Lorentzian (3,1) signature is the only minimal inner-product structure on the tangent spaces of M compatible with a single ordering direction, isotropic relational degrees of freedom, and a topologically invariant encoding of the ordering distinction. This is a plausibility derivation, not a rigorous theorem; it formalises the §12.2 (E2) motivation as a structural-selection argument but stops short of a measure-theoretic statement on the coarse-graining map. Nearby alternatives — degenerate metrics, structures with multiple ordering directions, non-metric tangent-space structures — are not strictly excluded; they are simply not minimal under the stated conditions.

Setup. Let $\mathcal{N}$ denote the substrate commitment network: a discrete structure consisting of commitment events and the ordering relation on them inherited from (C2). Let $\pi: \mathcal{N} \to M$ denote the coarse-graining map producing the smooth 4-manifold M (the existence and properties of π are content of (E1)). At each event $x \in M$, the tangent space T_xM acquires whatever inner product structure is inherited from the substrate-level statistics of commitment correlations along paths through $\pi^{-1}(x)$.

Argument. Pick an event $x \in M$ and a substrate event $y \in \pi^{-1}(x)$. The substrate-level ordering classifies any neighbouring substrate event y′ as one of three kinds with respect to y:

- *Future-directed:* y precedes y′ in the substrate ordering.
- *Past-directed:* y′ precedes y in the substrate ordering.
- *Order-comparable in neither direction:* y and y′ are not related by the ordering.

This classification is exhaustive (every pair of comparable or incomparable events falls into exactly one class) and the three classes are *structurally distinct*: future-directed and past-directed are inverses of each other under reversal of the ordering, while order-incomparable events are invariant under that reversal. The ordering therefore induces a substrate-level partition of neighbouring events into two disjoint kinds: *order-related* (future-directed or past-directed) and *order-unrelated*.

The coarse-graining map π pushes this partition forward to T_xM. Concretely, a tangent vector $v \in$ T_xM is realised as the equivalence class of substrate-level paths through $\pi^{-1}(x)$ that produce v in the coarse-graining limit. Each such path, near y, points either into the order-related or the order-unrelated class. The partition of substrate paths therefore induces a partition of tangent vectors:

- *Order-aligned tangent vectors*, realised by paths whose neighbouring events are order-related to y;
- *Order-orthogonal tangent vectors*, realised by paths whose neighbouring events are order-unrelated to y.

Why this points toward indefinite signature. The structural intuition is that the tangent-space inner product, being inherited from the substrate, must distinguish these two classes — and that the natural way for a symmetric bilinear form to distinguish a 1-dimensional sub-class (the ordering direction) from its complement (the order-orthogonal directions) is for it to assign opposite signs to the two.

We are explicit that this is a structural intuition, not a rigorous argument. In principle, a positive-definite (Riemannian) inner product on T_xM could distinguish the ordering direction by *magnitude* — for example, an anisotropic Riemannian metric in which the ordering direction has different norm from the spatial directions. Such a form would still treat all directions as "spacelike" in signature terms, while encoding the ordering through magnitude. What forces *indefinite signature* specifically is something stronger than mere distinguishability: it is the requirement that the ordering distinction be *topologically invariant* — preserved under smooth deformations, invariant under reversal of the ordering, and structurally coupled to the causal-evolution well-posedness of the field equations. Magnitude is rescalable under coordinate changes; sign is not.

The argument is therefore: *if* the substrate-level ordering distinction is to be a topologically invariant feature of the inner product (which is what is required for a globally consistent causal structure on M), then the inner product must encode the distinction via sign rather than magnitude, and the signature must be indefinite. Whether the substrate's coarse-graining map π actually produces a topologically-invariant rather than magnitude-encoded distinction is a question we cannot answer at the level of (E1)–(E2): it requires the explicit form of π, which is content of Paper 3.

Dynamical grounding for the topological-invariance assumption. The structural argument above can be supplemented by a physical one. Field equations on M must be consistent with the irreversibility content of (C2): they must admit a well-posed initial-value formulation in which past commitment data determine future commitment evolution but not vice versa. Mathematically, this requires *hyperbolic* field equations (admitting a Cauchy problem with a preferred direction of evolution) rather than elliptic ones (which are time-symmetric and have no notion of initial data evolving forward). Hyperbolicity of second-order linear partial differential equations on a Lorentzian manifold is precisely the property selected by *indefinite* signature: the principal symbol of the wave operator $g^{\mu\nu} \partial_\mu \partial_\nu$ is hyperbolic if and only if g has indefinite signature. A positive-definite (Riemannian) metric would give an elliptic principal symbol with no causal-evolution structure, regardless of whether it encodes a preferred direction by magnitude. The topological-invariance assumption (E2b) — that signature is sign-encoded rather than magnitude-encoded — is therefore not merely a structural choice but a *dynamical requirement* for the field equations on M to be consistent with irreversibility (C2). This grounds (E2b) in the construction's physics rather than in mathematical taste.

Given the topologically-invariant version, the signature has one timelike and three spacelike directions — one timelike because (C2) selects exactly one ordering direction, and three spacelike because M has dimension 4 (per (E1)).

What this argument establishes and what it does not. It establishes:

- (E2) is plausibly reducible to (E1) plus (C2)+(E4), modulo the topological-invariance assumption on the substrate's ordering distinction (named (E2b) in §12.2).

- Conditional on topological invariance, the number of timelike directions equals the number of independent substrate-level ordering directions, which is one by (C2).
- Conditional on topological invariance, the number of spacelike directions equals the manifold dimension minus the number of timelike directions: $4 - 1 = 3$.

It does not establish:

- That the coarse-graining map π actually exists and has the required properties; this is content of (E1) and is not derived in this paper.
- That π preserves the substrate-level ordering distinction *topologically* rather than via magnitude; this is the load-bearing intuition of the argument and is not proven.
- That the tangent-space inner product takes the specific form $g_{\mu\nu} = \mathrm{diag}(-1, +1, +1, +1)$ in any particular coordinate system; only the signature class is fixed.
- That the resulting signature is preserved at all events in M (i.e., that the manifold is globally Lorentzian rather than merely locally so); global signature continuity is a non-trivial topological assumption that we leave to (E1).

Status. This is a plausibility derivation with an identified gap. The structural intuition — substrate ordering selects one preferred direction, and a topologically-invariant encoding of that distinction in the inner product requires sign rather than magnitude — is a reasonable physical argument but not a theorem. The reduction of (E2) to (E1) plus (C2)+(E4) is real but conditional on the topological-invariance step (E2b), which awaits Paper 3.

For the purposes of the present paper, the upshot is that (E2) is not free: given (E1) and (E2b), Lorentzian signature is the minimal compatible structure, and (E2) can be understood as the natural selection within the inner-product structures admissible at the tangent-space level. The supplementary-inputs list could in principle be tightened by treating (E2) as derived from (E1)+(E2b), with (E2b) being the substantively independent component. We retain (E2) in the §12.2 list as the standard formulation for readability, with (E2b) named explicitly there as the genuine independent content. Conceptually, the count of *independent* supplementary inputs is now closer to three (E1, E3, E5) than five.

A.3 (E3) EFT power-counting from (C1) finite distinguishability

Claim. Standard effective-field-theory power-counting — the organisation of higher-dimension operators by powers of the cutoff scale $M_\star$ — follows from (C1) finite distinguishability via the Appelquist–Carazzone-type decoupling argument applied at the substrate cutoff scale $\ell_\star$. The dimensional structure (operators of mass dimension n appear with at least n powers of $M_\star^{-1}$) is genuinely derivable; the *naturalness* assumption (Wilson coefficients are O(1)) is not, and remains a residual input awaiting the substrate-level matching of Paper 3.

Argument: dimensional structure. (C1) asserts that the substrate admits a finite distinguishability scale $\ell_\star$, equivalently a UV scale $M_\star = \ell_\star^{-1}$ above which no observable distinguishes substrate states. In the post-emergence effective theory, this implies: any operator

$\mathcal{O}$ of mass dimension $d > 4$ appearing in the effective Lagrangian must enter with a coefficient $c_{\mathcal{O}}$ of mass dimension $4 - d$, hence

$$c_{\mathcal{O}} = \tilde{c}_{\mathcal{O}} \cdot M_\star^{(4-d)} ,$$

with $\tilde{c}_{\mathcal{O}}$ dimensionless. This follows from dimensional analysis given that the only UV scale in the theory is $M_\star$ from (C1) — there is no other dimensionful parameter that can appear in $c_{\mathcal{O}}$.

The argument that (C1) implies $M_\star$ is the *only* UV scale runs as follows. By (C1), no observable depends on sub-$\ell_\star$ structure, so no quantity in the effective theory can be sensitive to scales above $M_\star$ in any way other than through the cutoff itself. Any operator coefficient with explicit dependence on a scale $M \neq M_\star$ would either (i) introduce a new physical scale not present in (C1), violating minimality (C6), or (ii) be expressible as $M_\star$ times a dimensionless ratio $M/M_\star$, in which case $M_\star$ remains the organising scale.

The residual Wilson-coefficient gap. What the above argument does *not* establish is that $\tilde{c}_{\mathcal{O}} = O(1)$. It establishes only that $\tilde{c}_{\mathcal{O}}$ is dimensionless. A theory in which $\tilde{c}_{\mathcal{O}} \sim 10^6$ for some operator $\mathcal{O}$ is dimensionally consistent with (C1) but parametrically different from standard NDA expectations.

The standard EFT argument that $\tilde{c}_{\mathcal{O}} = O(1)$ requires either:

- (a) An explicit UV completion in which $\tilde{c}_{\mathcal{O}}$ is computed by matching;
- (b) A naturalness assumption that no parametrically large dimensionless numbers appear without physical reason;
- (c) A symmetry argument that forbids the operator from appearing with anomalously large coefficient.

None of (a)–(c) is supplied by (C1) alone. The substrate-level coarse-graining map (deferred to Paper 3) would, in principle, produce explicit $\tilde{c}_{\mathcal{O}}$ by matching — but until that calculation is done, the $O(1)$ assumption is an additional input.

What (E3) reduces to. Given the above, (E3) splits into two sub-inputs:

- *(E3a) Dimensional cutoff structure.* Operators of dimension $d > 4$ enter the effective Lagrangian with coefficient $\sim M_\star^{(4-d)}$. This is *derivable* from (C1) via dimensional analysis and minimality (C6). It is not an independent input.
- *(E3b) Wilson coefficient naturalness.* The dimensionless coefficients $\tilde{c}_{\mathcal{O}}$ are $O(1)$. This is *not derivable* from (C1)–(C6) and remains a supplementary input. Rigorous derivation requires the substrate-level matching of Paper 3.

The §12.2 (E3) entry conflates these two pieces. A future revision might split (E3) into (E3a) and (E3b), with (E3a) folded into the closure axioms and only (E3b) retained as a supplementary input.

Status. Partial derivation, with explicit residual gap. The dimensional half of (E3) is a genuine consequence of (C1)+(C6); the naturalness half awaits Paper 3.

A.4 Summary

The appendix has shown:

- **(E4) is contained in (C2) at the level of existence.** The substrate-level ordering whose existence (E4) asserts is presupposed by (C2)'s notion of irreversibility; the structural properties of that ordering remain tied to (E1).
- **(E2) reduces to (E1) plus (C2), conditional on (E2b) topological-invariance of the ordering distinction.** Lorentzian signature is the minimal inner-product structure compatible with a single ordering direction, isotropic relational degrees of freedom, and a topologically invariant encoding of the ordering distinction. (E2b) is itself substrate-level and grounded dynamically in A.2 by the requirement that field equations on M be hyperbolic.
- **(E3) splits into (E3a) and (E3b).** Dimensional cutoff structure follows from (C1)+(C6); Wilson-coefficient naturalness is a residual input awaiting Paper 3.

The two inputs not addressed are:

- **(E1)** — derivation of the smooth 4-manifold from substrate dynamics. Belongs to the K=7-closure and time-emergence sectors of the corpus; not attempted here.
- **(E5)** — pointwise commitment commutativity. A derivation from (C1)–(C5) is not available within the present construction, and (E5) is therefore taken as an independent supplementary input. Re-opening here would either rediscover the (C2)+(C3) framing whose limitations are discussed in §9.1, or require importing substrate-algebra machinery beyond the paper's scope.

Net effect on the supplementary-inputs count. What was presented in the main text as five independent supplementary inputs reduces, on closer inspection, to:

- **One input contained at the level of existence within a closure axiom:** (E4) within (C2), with structural properties tied to (E1).
- **One input conditionally reducible to another:** (E2) to (E1) + (C2), modulo (E2b) — the topological-invariance condition of A.2.
- **One input partially derivable, partially residual:** (E3) split into (E3a, derived) and (E3b, residual).
- **Genuinely independent inputs:** (E1) and (E5) — plus (E3b), and conditionally (E2b) from (E2)'s reduction — which remain the substantive supplementary content.

The number of independent structural inputs can therefore be reduced to approximately three: (E1), (E3b), and (E5), with (E2) reducible to (E1) modulo (E2b) — the substrate-level topological-invariance condition deferred to Paper 3 and grounded dynamically in A.2 by the irreversibility requirement on field equations. The exact count depends on accounting choices — whether (E2b) is treated as substantive (raising the count to four) or absorbed under (E2)'s deferral to Paper 3 (keeping it at three), whether (E3a) and (E3b) are split or kept together — but the substantive content of the supplementary inputs effectively reduces to three core inputs plus the (E2b) conditional. The other appearances of "supplementary input" in the main text are bookkeeping conveniences for readers tracking the dependency structure, not statements of independent assumption.

We do not propose to revise the main body's framing. The five-input presentation is more readable and the dependency table in §12 is doing real work as a teaching device. But for a reader concerned about minimality of the assumption set, the appendix establishes that the construction's *substantive* dependencies beyond the closure axioms are: a smooth 4-manifold (E1), the naturalness of effective-Lagrangian Wilson coefficients (E3b), and pointwise commitment commutativity (E5). The other entries in the §12.2 list are either definitional or reducible to these together with (C1)–(C6).

A.5 What remains for Paper 3

The appendix above is, in part, a sharpened statement of the work Paper 3 needs to do. Specifically, Paper 3 must:

1. Construct the substrate-level commitment network $\mathcal{N}$ and the coarse-graining map $\pi: \mathcal{N} \to M$ producing the smooth 4-manifold (closing the residual content of (E1)) and verify that π preserves the substrate-level ordering distinction topologically (closing (E2b) and thereby completing the reduction of (E2) to (E1)).
2. Compute Wilson coefficients $\tilde{c}_\mathcal{O}$ by matching to the substrate dynamics (closing (E3b)).
3. Derive pointwise commitment commutativity from the substrate-level commitment algebra, or alternatively identify the substrate-level conditions under which commutativity holds and characterise the corrections when it fails (closing (E5)).

Items (1)–(3) define the scope of Paper 3 in the master-action roadmap. The present appendix does not attempt them, but it identifies them as the irreducible substrate-level inputs the rest of the corpus must supply. The main paper's claim — that gravity emerges as the unique tensorial closure of record dynamics under (C1)–(C6) plus (E1)–(E5) — is unchanged; what the appendix adds is that the *real* irreducible content of (E1)–(E5) is contained in three sub-inputs (the manifold itself, Wilson-coefficient naturalness, and pointwise commutativity), with the remainder being either definitional or structurally derivable.

Paper 1 of the four-paper VERSF Master Action roadmap. Sequel papers address (Paper 2) the dimensional bookkeeping and TPB[Φ] ~ $|\Phi|^2$ identity in detail; (Paper 3) the coarse-graining map producing $g_{\mu\nu}$ from $\Phi_{\mu\nu}$; (Paper 4) numerical determination of a, b, κ from the Standard Model matter sector.