

Article

Not peer-reviewed version

Guaranteed Tensor Luminality from Symmetry: A PT-Even Palatini Torsion Framework

[Chien-Chih Chen](#) *

Posted Date: 2 January 2026

doi: 10.20944/preprints202509.2421.v4

Keywords: modified gravity; Palatini formalism; torsion; PT symmetry; gravitational waves; luminal propagation



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Guaranteed Tensor Luminality from Symmetry: A PT-Even Palatini Torsion Framework

Chien-Chih Chen 

Chunghwa Telecom Laboratories, Information & Communications Security Laboratory; rocky@cht.com.tw

Abstract

Multimessenger constraints tightly bound the gravitational-wave speed to be luminal, posing a strong filter for modified gravity. This paper develops a symmetry-selected Palatini framework with torsion in which exact luminality at quadratic order is achieved by construction rather than by parameter tuning. Two ingredients shape the observable sector: (i) a scalar PT projector that keeps scalar densities real and parity-even, and (ii) projective invariance implemented via a non-dynamical Stueckelberg compensator that enters only through its gradient. Under an explicit posture (A1–A6), we establish three structural results: (C1) algebraic uniqueness of torsion to a pure-trace form aligned with the compensator gradient; (C2) bulk equivalence, modulo improvements, among a rank-one determinant route, a closed-metric deformation, and a PT-even CS/Nieh–Yan route; and (C3) a coefficient-locking identity that enforces $K = G$ for tensor modes on admissible domains, hence $c_T = 1$ with two propagating polarizations. Beyond leading order, the framework yields a distinctive, falsifiable next-to-leading correction $\delta c_T^2(k) = b k^2 / \Lambda^2$ (for $k \ll \Lambda$), predicting slope 2 in log–log fits across frequency bands (PTA/LISA/LVK). The analysis is formulated to be reproducible, with a public repository providing figure generators, coefficients, and tests that directly validate (C1)–(C3).

Keywords: modified gravity; Palatini formalism; torsion; PT symmetry; gravitational waves; luminal propagation

1. Introduction

Positioning & Scope. *This section motivates the problem and clarifies the paper’s scope for the present framework: we aim at a symmetry-based, conditionally luminal tensor sector at quadratic order, not at a full replacement of Λ CDM or a comprehensive cosmological fit. Our contribution is foundational: it identifies a structural route to $c_T = 1$ under explicit assumptions (A1–A6), and it frames a clear, falsifiable NLO prediction accessible to multi-band GW data.*

The dawn of multimessenger astronomy, heralded by the joint observation of gravitational waves (GW170817) and a gamma-ray burst (GRB170817A), has transformed fundamental physics [28,29]. The near-simultaneous arrival of these signals established that gravitational waves propagate at the speed of light to an astonishing precision, $|c_T/c - 1| \lesssim 10^{-15}$. This single measurement acts as a stringent filter for a vast landscape of modified gravity theories. While some models can be salvaged by fine-tuning parameters, such an approach is often considered theoretically unnatural. This predicament sharpens a fundamental question: *Can we construct a gravitational theory where luminal tensor speed is a direct consequence of an underlying symmetry principle, rather than an accident of parameter choice?*

This paper provides an affirmative answer within a first-order, Palatini-type geometry endowed with torsion. We propose that the observable sector of gravity is governed by two key symmetry principles:

1. **A Scalar PT Projection:** We postulate that all observable scalar densities are projected into a PT-even (parity-time symmetric) subspace. This acts as a selection rule, ensuring that the action is real and the Hamiltonian is Hermitian, while systematically eliminating potential parity-odd effects like gravitational birefringence at leading order.
2. **Projective Invariance:** We implement projective symmetry via a non-dynamical Stueckelberg compensator field $\epsilon(x)$, which can be intuitively understood as a geometric phase. This ensures that observables are only sensitive to the projectively invariant trace of torsion, $\mathcal{T}_\mu \equiv T_\mu - \partial_\mu \epsilon$, structurally avoiding extra propagating degrees of freedom that often plague torsionful theories.

Within this symmetric and self-consistent posture, formalized by six assumptions (A1–A6) in Section 2, we derive a cascade of powerful, conditional results at the quadratic action level. As illustrated in Figure 1, these principles first act as a structural constraint that aligns the torsion into a unique pure-trace form (Result C1). This simplification then reveals a remarkable bulk equivalence between three distinct theoretical routes (Result C2).

The culmination of this structure is a novel *coefficient locking* mechanism (Result C3). Demanding the absence of unphysical mixing between tensor modes and non-propagating fields fixes a unique relationship between the theory’s building blocks. This locking forces the kinetic (K) and gradient (G) terms in the tensor action to be identical up to a boundary term. This identity, $K = G$, structurally ensures an *exactly luminal tensor speed* ($c_T = 1$) by construction. This result is algebraic, requires no parameter tuning, and emerges directly from the imposed symmetries.

The framework’s utility lies not only in its structural elegance but also in its predictive power. It yields a distinctive, testable signature at the next-to-leading order (NLO) that is absent in standard General Relativity: a frequency-dependent deviation in the speed of gravity given by

$$\delta c_T^2(k) \equiv c_T^2(k) - 1 = b \frac{k^2}{\Lambda^2} \quad (k \ll \Lambda), \quad (1)$$

where k is the wavenumber and Λ is the effective energy scale of new physics. A detection of this specific k^2 scaling in future multi-band GW data (from Pulsar Timing Arrays, LISA, and next-generation detectors) would provide a powerful probe of spacetime’s fundamental geometry, while its absence would falsify the framework. This paper lays out the theoretical foundations, proves the core properties, and details the clear, falsifiable predictions of this symmetry-based approach.




Figure 1. Roadmap of the quadratic analysis (bulk-only). The scalar PT projector (Thm. 1) defines the allowed observable sector. Projective invariance is realized via $\mathcal{T}_\mu = T_\mu - \partial_\mu \epsilon$. (C1) The posture enforces a pure-trace torsion aligned with the spurion gradient (Section 4, Equation (11)). (C2) It establishes the bulk equivalence of three construction routes up to boundary terms (Section 5, Equation (23)). (C3) A locking condition then yields $K = G$ and thus $c_T = 1$ by construction (Section 6, Equation (34)).

2. The Theoretical Posture: Symmetries, Assumptions, and Projectors

Positioning & Scope. We specify the working posture and its admissible variational domains. The two-derivative truncation and the spurion limit are deliberate, conservative choices ensuring a real, PT -even scalar sector while preserving projective invariance in observables. All claims in later sections are to be read strictly under assumptions A1–A6 and on admissible boundaries; outside this posture no general statements are implied.

This section fixes the theoretical foundation of our framework. We begin by stating the two guiding principles already previewed in Section 1: (i) projective invariance in observables, implemented via a Stueckelberg compensator, and (ii) a scalar PT projection that enforces reality at the level

of observable densities. We then formalize the complete set of assumptions (A1–A6) defining our working posture, and finally define the scalar PT projector and its key properties. All constructions are performed within a metric-compatible Palatini formalism on an oriented $(3+1)$ -dimensional spacetime, with torsion $T^A \equiv de^A + \omega^A_B \wedge e^B$ and curvature $R^{AB} \equiv d\omega^{AB} + \omega^A_C \wedge \omega^{CB}$.

2.1. Principle 1: Projective Symmetry and the Stueckelberg Spurion

Our first guiding principle is projective symmetry, a standard redundancy of Palatini geometry under which the affine connection shifts without changing the curvature tensor:

$$\Gamma^\alpha_{\mu\nu} \rightarrow \Gamma^\alpha_{\mu\nu} + \delta^\alpha_\mu \xi_\nu, \quad T_\mu \equiv T^\alpha_{\mu\alpha} \rightarrow T_\mu + (d-1) \xi_\mu. \quad (2)$$

To ensure physical observables are invariant under (2), we introduce a Stueckelberg compensator field $\epsilon(x)$ and assign it the compensating transformation $\partial_\mu \epsilon \rightarrow \partial_\mu \epsilon + (d-1) \xi_\mu$. This allows the theory to be organized in terms of the *projectively invariant* torsion-trace combination

$$\boxed{\mathcal{T}_\mu \equiv T_\mu - \partial_\mu \epsilon} \quad (\text{projectively invariant}). \quad (3)$$

Within our two-derivative truncation we adopt a *spurion limit*: ϵ is treated as a non-dynamical background field at this order. Consequently, admissible observables may depend on $\partial\epsilon$ only through the invariant combination $\mathcal{T}_\mu$. This conservative choice preserves projective invariance while enabling the algebraic simplifications central to our results.

2.2. Optional UV Motivation: A PTQ Interpretation (Not a Premise)

Positioning & Scope. This subsection is interpretive and optional. It provides one natural UV narrative for the compensator $\epsilon(x)$ and for the scalar- PT projection, but none of the proofs in this paper depend on that narrative. All structural results (C1–C3) rely only on the spurion/projective-invariant chain under the posture A1–A6.

Why include an optional UV interpretation?

The posture adopted here—a non-dynamical Stueckelberg compensator entering only through $\mathcal{T}_\mu \equiv T_\mu - \partial_\mu \epsilon$ and a scalar- PT projection acting on observable densities—is fully consistent as a conservative two-derivative effective framework. Nevertheless, it is conceptually useful to note that the same pattern can arise as a particularly economical low-energy remnant of a richer UV completion. One such UV viewpoint is a PT -symmetric quaternionic-spacetime (PTQ) interpretation, in which ϵ is naturally read as a phase/clock-like geometric spurion and the scalar- PT projection acts as a reality filter on the observable sector. We emphasize that this PTQ narrative is a *motivation* and a *roadmap* for deepening the theory, not an additional assumption.¹

Minimal mapping (interpretive only).

From a PTQ perspective, the two guiding principles may be viewed as IR footprints of a UV structure:

1. **Projective invariance as an IR gauge redundancy.** A projective shift of the Palatini connection may be interpreted as an emergent redundancy after integrating out heavy UV modes, leaving only the invariant combination $\mathcal{T}_\mu$ in observables. In this reading, ϵ is the minimal Stueckelberg spurion that packages UV phase information into the IR without introducing additional propagating degrees of freedom at quadratic order.
2. **Scalar- PT projection as a selection rule on observables.** A PTQ-inspired completion naturally motivates imposing a reality condition at the level of *scalar densities* (rather than on all microscopic

¹ Readers may skip this subsection without losing any logical step of the main argument. In particular, no UV-specific input is used in establishing C1–C3.

fields). Operationally, this is exactly what Π_{PT} implements in our posture: it retains the PT -even scalar sector in observables while keeping the underlying geometric variables (e, ω) otherwise standard within metric-compatible Palatini gravity.

What depends on PTQ (and what does not).

To prevent over-reading, we explicitly separate *structure* from *interpretation*:

- **Independent of PTQ:** The operator basis, Palatini algebraicity, the uniqueness map (C1), the three-route bulk equivalence (C2), and the coefficient-locking identity enforcing $c_T = 1$ (C3) follow solely from A1–A6 together with the spurion/projective-invariant construction and standard variational identities.
- **PTQ as a UV guide (optional):** PTQ provides a plausible UV interpretation for why a non-dynamical phase-like spurion and a scalar-level PT selection may be stable low-energy remnants, and it offers an intuitive reading of the EFT cutoff Λ entering NLO dispersion as a scale at which additional UV degrees of freedom would become dynamical.

Practical takeaway.

In short, PTQ is the *more ambitious* version of the story—it suggests a concrete UV direction and a principled origin for (ϵ, Π_{PT}) —but the present paper deliberately treats it as *non-essential*. The main claims are symmetry-protected and auditable within the spurion/projective-invariant posture, regardless of which UV completion one ultimately adopts (PTQ or otherwise).

2.3. Principle 2: The Scalar PT Projector and Reality of Observables

Our second guiding principle is to project observable scalar densities onto a real, PT -even subspace. Operationally, this acts as a “reality filter” at the level of *observables*: the projected action is real by construction, and the quadratic operators governing physical perturbations are Hermitian within the posture specified below.

Why PT on scalar densities?

On an oriented Lorentzian manifold, the combined PT transformation preserves the orientation of the spacetime volume element (A2). Imposing the reality condition at the level of *scalar densities* is the minimal intervention compatible with the Palatini variational principle: it does not require imposing PT on every microscopic field component, yet it ensures that all scalar observables entering the action are PT -even in the projected sector.

Connection to Hermiticity (within the posture).

We impose the projection on observable densities rather than as a microscopic constraint on (e, ω) . Within our admissible variational domains and with the spurion held fixed, the projection commutes with the Palatini variation (A3), so the resulting quadratic forms in the physical sector are Hermitian:

“ PT -even projected scalar densities $\Rightarrow$ real action density $\Rightarrow$ Hermitian quadratic operators (in the projected sector).”

This statement is to be read strictly under A1–A6.

2.4. The Working Posture (Assumptions A1–A6)

The conditional results of this paper are derived under the following explicit assumptions, collectively referred to as our *posture*.

Assumptions A1–A6 (The working posture).	
A1.	<i>PT-invariant domain/measure.</i> For any scalar density X , the integration domain is such that $\int X = \int \mathcal{PT}[X]$.
A2.	<i>Orientation and Hodge dual.</i> The combined PT transformation preserves spacetime orientation and commutes with the Hodge star operator, $[\mathcal{PT}, *] = 0$.
A3.	<i>Projection-variation commutation (scalar densities).</i> For scalar densities $\mathcal{O}$, the operations of variation and projection commute: $\delta \Pi_{\text{PT}}[\mathcal{O}] = \Pi_{\text{PT}}[\delta \mathcal{O}]$ (proved as Lemma 1 under the admissibility conditions below).
A4.	<i>Admissible boundaries/falloffs.</i> The variational domain is restricted to spacetimes where improvement currents are pure boundary terms that do not contribute to the symplectic form (e.g., compact domains or standard asymptotically flat/FRW fall-offs).
A5.	<i>Nieh–Yan as a boundary term.</i> The Nieh–Yan 4-form $\text{NY} \equiv d(e^A \wedge T_A)$ is treated as a total derivative affecting only boundary terms.
A6.	<i>Two-derivative truncation and spurion limit.</i> The analysis is restricted to quadratic order in fields with at most one derivative per building block in the tensor sector. The compensator ϵ is treated as a non-dynamical background field (spurion) at this order.

Note on admissible domains.

The posture admits standard, simply connected spacetimes. Manifolds with torsion defects, nontrivial holonomy for ϵ , or multi-valued spurions could induce non-vanishing boundary fluxes and may spoil the bulk equivalences derived later (cf. Section 5).

2.5. Formal Properties of the PT Projector

We now formalize the scalar PT projector and its essential properties.

Convention note (spurion charges).

We treat ϵ as a non-dynamical spurion field. Discrete P and T assignments for ϵ are chosen so that the invariant combination $\mathcal{T}_\mu = T_\mu - \partial_\mu \epsilon$ is compatible with the projected PT -even observable sector. This ensures that PT selection rules act consistently on the spurion-built scalars used below.

Action of P , T , and PT .

The transformations of key geometric objects under P and T are summarized in Table 1. We emphasize that T is anti-unitary: it includes complex conjugation on scalar coefficients.

Table 1. Action of P and T on representative geometric objects. T is anti-unitary.

Object	P	T	Note
e^A_μ (vierbein)	+	+	
ω^{AB}_μ (connection)	+	−	Anti-unitary action of T
$T^A_{\mu\nu}$ (torsion)	+	−	
$R^{AB}_{\mu\nu}$ (curvature)	+	+	
ϵ (spurion field)	+	−	Chosen spurion assignment
$\partial_\mu \epsilon$ (gradient)	−	+	Vector-like under P
$\epsilon^{\mu\nu\rho\sigma}$ (density)	−	−	PT-even

Definition and core properties.

The projector onto the PT -even subspace of scalar densities is defined by

$$\Pi_{\text{PT}}[\mathcal{O}] \equiv \frac{1}{2}(\mathcal{O} + \mathcal{PT}[\mathcal{O}]), \quad \Pi_{\text{PT}}^2 = \Pi_{\text{PT}}.$$

(4)

Under assumption A1 (a PT -invariant integration domain/measure), this projection is self-adjoint in the sense appropriate for scalar densities on the admissible domain, and it defines a consistent restriction to a PT -even observable sector.

Lemma 1 (Projection-Variation Commutation (A3)). *Under A1 and A4, if the spurion ϵ is held fixed (non-variational), then for variations of (e, ω) one has $\delta \Pi_{PT}[\mathcal{O}] = \Pi_{PT}[\delta \mathcal{O}]$ for any scalar density $\mathcal{O}$.*

Proof. Linearity gives $\delta \Pi_{PT}[\mathcal{O}] = \frac{1}{2}(\delta \mathcal{O} + \delta \mathcal{PT}[\mathcal{O}])$. On the admissible domain (A4), boundary improvements do not contribute to the symplectic form, and with ϵ held fixed the PT map acts only on the dynamical variations of (e, ω) and on complex conjugation (anti-unitarity of T). Therefore $\delta \mathcal{PT}[\mathcal{O}] = \mathcal{PT}[\delta \mathcal{O}]$, implying the claim. $\square$

Theorem 1 (PT Selection Rules). *Under the posture A1–A5, for any scalar density $\mathcal{O}$: (i) if $\mathcal{PT}[\mathcal{O}] = +\mathcal{O}$, then $\Pi_{PT}[\mathcal{O}] = \mathcal{O}$; (ii) if $\mathcal{PT}[\mathcal{O}] = -\mathcal{O}$, then $\Pi_{PT}[\mathcal{O}] = 0$.*

Proof. Immediate from the definition (4). $\square$

2.6. Summary of Key Ingredients

For clarity, we consolidate the key spurion-built variables that constitute the observable scalar sector in this framework.

Key Ingredients for the Observable Sector

- *Projected spurion scalar:* $\Sigma_\epsilon \equiv \Pi_{PT}[(\partial\epsilon)^2]$. The fundamental real, PT -even scalar built from the spurion gradient.
- *Projectively invariant trace:* $\mathcal{T}_\mu \equiv T_\mu - \partial_\mu \epsilon$. The only combination through which the torsion trace enters observables.
- *Quadratic torsion invariant:* $I_T \equiv -\Pi_{PT}[T^A_{BC} T_A^{BC}]$. The primary quadratic invariant built from the full torsion tensor.

Note on regularity: On loci where $\Sigma_\epsilon = 0$, related quantities (e.g., $\hat{n}_\mu$) are defined by a smooth continuation limit.

3. Building the Operator Basis: From Symmetries to Observables

Positioning & Scope. *This section is technical by design and self-contained. It enumerates the complete quadratic, PT -even, projectively invariant basis before equations of motion (pre-lock), then shows its collapse after solving for the torsion (post-lock). The goal is transparency: all later identities are traceable to this basis and to integration-by-parts on admissible domains.*

Having established our guiding principles and working posture, we now construct the complete basis of observable operators at the lowest non-trivial order (quadratic in fields, with at most one derivative per field). This systematic process is a crucial preparatory step. We first identify all operators allowed by the symmetries *before* applying the equations of motion, defining the "pre-lock" basis. We then show how this basis simplifies once the dynamics are solved in Section 4, yielding the "post-lock" sector. This ensures our theory is built from a complete and self-consistent set of building blocks.

Throughout this section, the posture A1–A6 defined in Section 2 is in force. All scalar densities are implicitly understood to be projected via Π_{PT} , rendering them real and PT -even.

3.1. Symmetry-Allowed Building Blocks

The combined requirements of projective invariance and PT symmetry act as a powerful filter. As established in Section 2, all observables must be constructed from the projectively invariant trace of torsion, $\mathcal{T}_\mu \equiv T_\mu - \partial_\mu \epsilon$, and the Stueckelberg gradient $\partial_\mu \epsilon$ itself.

At quadratic order with at most one derivative per field, Theorem 1 allows only three types of PT -even scalar monomials to survive the projection:

- (i) The quadratic torsion invariant, built from $T^A{}_{BC} T_A{}^{BC}$.
- (ii) The quadratic spurion-gradient invariant, $(\partial\epsilon)^2$.
- (iii) The mixing term, $T_\mu \partial^\mu \epsilon$.

Crucially, before any equations of motion are applied, the mixing term is an independent operator.² This leads to a two-stage structure for our operator basis.

3.2. Two-Stage Basis Closure: Pre-Lock and Post-Lock

The analysis of the operator basis naturally separates into two stages, distinguished by whether the dynamical solution for torsion (Result C1, proven in Section 4) has been utilized.

The Pre-Lock Basis.

Before solving the equations of motion, the three surviving monomials form a complete set of independent operators. We denote this basis as:

$$\text{Pre-lock basis: } I_1 \equiv \Pi_{\text{PT}}[T^A{}_{BC} T_A{}^{BC}], \quad I_2 \equiv \Sigma_\epsilon \equiv \Pi_{\text{PT}}[(\partial\epsilon)^2], \quad I_3 \equiv \Pi_{\text{PT}}[T_\mu \partial^\mu \epsilon]. \quad (5)$$

The Post-Lock Sector.

After solving the Palatini equations of motion (Section 4), we obtain the algebraic alignment $T_\mu = 3\eta \partial_\mu \epsilon$. This condition is a dynamical result derived from the variation of the pre-lock action. Substituting this solution back into the basis operators causes a collapse: the mixing term I_3 becomes proportional to I_2 .

$$I_3 \xrightarrow{\text{C1: } T_\mu = 3\eta \partial_\mu \epsilon} \Pi_{\text{PT}}[(3\eta \partial_\mu \epsilon) \partial^\mu \epsilon] = 3\eta \Sigma_\epsilon = 3\eta I_2. \quad (6)$$

Furthermore, the torsion invariant I_1 also becomes proportional to I_2 . To maintain consistency with standard conventions, we define the *normalized* torsion invariant I_T and describe the post-lock physics as follows:

$$I_T \equiv -I_1 = -\Pi_{\text{PT}}[T^A{}_{BC} T_A{}^{BC}]. \quad (7)$$

Post-lock sector: The physical sector becomes one-dimensional, spanned by Σ_ϵ (or equivalently I_T). (8)

The following lemma formalizes this two-stage structure, with its proof relying on standard tensor decomposition and the self-adjointness of the projector under the A1–A5 posture.

Lemma 2 (Two-stage closure). *Under assumptions A1–A5, any PT-even quadratic scalar density $\mathcal{O}_{(2)}(T, \partial\epsilon)$ with at most one derivative per factor can be expanded on the pre-lock basis $\{I_1, I_2, I_3\}$ up to a total divergence. After solving for the torsion (C1), the observable sector collapses to a single independent invariant.*

3.3. The Action Skeleton and Order of Operations

The closure of the operator basis allows us to write down the most general action at this order. The minimal bulk action skeleton is first written in the pre-lock basis:

$$S_{\text{bulk}} = \int d^4x \sqrt{-g} \left[\frac{M_{\text{Pl}}^2}{2} R(e, \omega) + \alpha_1 I_1 + \alpha_2 I_2 + \alpha_3 I_3 \right], \quad (9)$$

where $\alpha_{1,2,3}$ are free real coefficients. It is the variation of *this* action (specifically the terms $\alpha_1 I_1 + \alpha_3 I_3$) that generates the linear equations of motion for torsion, leading to the solution $T_\mu \propto \partial_\mu \epsilon$.

After integrating out the torsion (i.e., substituting the solution back), the effective action depends on a single parameter combination:

$$S_{\text{bulk}}^{(\text{post-lock})} = \int d^4x \sqrt{-g} \left[\frac{M_{\text{Pl}}^2}{2} R(e, \omega) + (\alpha_2 + 3\eta \alpha_3 - 6\eta^2 \alpha_1) \Sigma_\epsilon \right]. \quad (10)$$

² Before a dynamical relation between T_μ and $\partial_\mu \epsilon$ is established, no integration-by-parts identity can reduce the term $T_\mu \partial^\mu \epsilon$ to a combination of T^2 and $(\partial\epsilon)^2$ up to boundary terms.

As proven in Section 4, the dynamics impose the crucial relation $I_T = -6\eta^2 \Sigma_\epsilon$, confirming that the post-lock physics is fully described by a single independent degree of freedom at this order.

Order of operations and consistency.

This two-stage structure defines a clear and consistent logical pipeline for our analysis. The scalar projector Π_{PT} acts at the level of densities, and by Lemma 1, we may *project then vary*. The alignment of torsion is an *algebraic solution* to the equations of motion; it is not a projection. The correct procedure is therefore:

- (i) Enumerate & project $\Rightarrow$ Pre-lock basis $\{I_1, I_2, I_3\}$
 $\xrightarrow{\text{vary } \omega}$ (ii) Solve algebraic EOM $\Rightarrow$ Post-lock sector $\{I_T \propto \Sigma_\epsilon\}$.

This path ensures that all dynamical consequences are derived consistently from the foundational symmetries.

4. The First Pillar: Uniqueness of Torsion (Result C1)

Positioning & Scope. *Here we show an algebraic result—pure-trace torsion aligned with the spurion gradient—within the stated posture. The statement is limited to quadratic order and to admissible domains; it does not claim non-perturbative uniqueness. The immediate function of C1 is to reduce the operator content and to remove axial/traceless channels from the observable sector at this order.*

This section presents the first major result (C1) of our framework. Having prepared the operator basis in Section 3, we now solve the Palatini equations of motion within our symmetric posture. The combination of the *PT* projector and projective invariance acts as a powerful selection principle, uniquely determining the algebraic form of the torsion tensor. We will show that of all possible forms of torsion (trace, axial, and tensor parts), only a pure-trace component aligned with the Stueckelberg gradient can exist in the observable sector. This result is not an assumption but a direct consequence of the symmetries. It dramatically simplifies the theory by collapsing the operator basis and paves the way for the equivalences and luminosity proof in the subsequent sections.

The central theorem of this section is as follows:

Theorem 2 (Palatini–*PT* Uniqueness of Torsion (C1)). *Within the A1–A6 posture, the Palatini equations of motion fix the algebraic torsion to be uniquely of the pure-trace form*

$$T^A{}_{BC} = 2\eta \delta^A{}_{[B} \partial_{C]} \epsilon, \quad \eta > 0. \quad (11)$$

This implies that the axial (S^μ) and traceless tensor ($q_{\lambda\mu\nu}$) irreducible representations of torsion must vanish. A direct consequence is the key algebraic identity relating the operators in the post-lock basis:

$$I_T = -6\eta^2 \Sigma_\epsilon. \quad (12)$$

This identity demonstrates that, dynamically, the two operators in our post-lock basis are not independent. The remainder of this section is dedicated to proving this theorem and exploring its implications.




Figure 2. Irreducible torsion content at quadratic order (log-log view). Ratios of projected scalar strengths comparing the pure-trace block against the axial and traceless blocks. The Palatini algebraicity (Section 4.1) and the PT projector dynamically drive the axial and traceless components to zero, leaving the pure-trace map (11) as the unique survivor. [nb: fig_c1_pure_trace.py]

$$\omega = \dot{\omega} + K(T)$$



Figure 3. Connection decomposition and the C1 map. Top: The Levi-Civita/contorsion split $\omega = \dot{\omega} + K(T)$. Bottom: The three irreducible representations of torsion $(T_\mu, S^\mu, q_{\lambda\mu\nu})$ are subjected to the C1 alignment, which enforces $S^\mu = 0$, $q_{\lambda\mu\nu} = 0$, and $T_\mu = 3\eta \partial_\mu \epsilon$ (11). This dynamically implies the relation $I_T = -6\eta^2 \Sigma_\epsilon$ (12).

4.1. Proof of Theorem 2

Our proof proceeds in three logical steps, which we now detail.

4.1.1. Step 1: Construct the Most General Ansatz for Torsion

At the one-derivative level, the most general Lorentz-covariant linear ansatz for a torsion tensor built from a single covector $\partial_\mu \epsilon$ can be written as

$$T^A_{BC} = A \delta^A_{[B} \partial_{C]} \epsilon + B \epsilon^A_{BCD} \partial^D \epsilon + (\mathcal{P}_q \cdot \partial \epsilon)^A_{BC}, \quad (13)$$

where A and B are real coefficients, and the three terms correspond to the trace, axial, and traceless tensor parts, respectively. However, a key result from group theory (proven in Appendix B, Proposition B.1) is that it is impossible to construct a non-zero traceless, rank-3 tensor $(q_{\lambda\mu\nu})$ linearly from a single covector. Therefore, the third term vanishes identically, $(\mathcal{P}_q \cdot \partial \epsilon)^A_{BC} = 0$. Our ansatz simplifies to

$$T^A_{BC} = A \delta^A_{[B} \partial_{C]} \epsilon + B \epsilon^A_{BCD} \partial^D \epsilon. \quad (14)$$

While the PT projector eliminates certain parity-odd scalar combinations, it does not by itself remove the axial part (the B term). The Palatini equations of motion are required to achieve this.

4.1.2. Step 2: Solve the Palatini Equations of Motion

We now determine the coefficients A and B by varying the pre-lock bulk action (9) with respect to the connection. This action is built from the linearly independent basis operators $\{I_1, I_2, I_3\}$ defined in Equation (5). Specifically, the torsion-dependent part is:

$$S_{\text{torsion}} = \int d^4x \sqrt{-g} [\alpha_1 I_1 + \alpha_3 I_3] = \int d^4x \sqrt{-g} [\alpha_1 \Pi_{\text{PT}}(T^2) + \alpha_3 \Pi_{\text{PT}}(T_\mu \partial^\mu \epsilon)]. \quad (15)$$

By Lemma 1, we can vary the total action with respect to the spin connection ω *after* projection. The variation yields independent algebraic equations for the three torsion irreps, thanks to the block-diagonal structure of the kinetic term (established in Lemma 3).

Lemma 3 (Block-diagonality of variation). *In metric-compatible Palatini gravity, the linear map from a variation in the connection, $\delta\omega^{AB}{}_\mu$, to the variations of the three torsion irreps, $(\delta T_\mu, \delta S_\mu, \delta q_{\lambda\mu\nu})$, is blockwise non-degenerate. (Proof: see explicit decomposition in Appendix B).*

Using the explicit coefficients for T^2 in the irrep basis (derived in Appendix B), the Euler–Lagrange equations for each irrep are:

$$\text{(vector)} \quad \frac{4\alpha_1}{3} T_\mu + \alpha_3 \partial_\mu \epsilon = 0 \quad \Rightarrow \quad T_\mu = -\frac{3\alpha_3}{4\alpha_1} \partial_\mu \epsilon \equiv 3\eta \partial_\mu \epsilon, \quad (16)$$

$$\text{(axial)} \quad \frac{\alpha_1}{12} S_\mu = 0 \quad \Rightarrow \quad \boxed{S^\mu = 0}, \quad (17)$$

$$\text{(traceless)} \quad 2\alpha_1 q_{\lambda\mu\nu} = 0 \quad \Rightarrow \quad \boxed{q_{\lambda\mu\nu} = 0}. \quad (18)$$

Here we have defined the parameter $\eta \equiv -\alpha_3/(4\alpha_1)$. Assuming $\alpha_1 \neq 0$ (a non-trivial torsion sector), the axial and traceless equations immediately force their corresponding torsion components to vanish. This dynamically sets the coefficient $B = 0$ in our ansatz (14). The vector equation (16) uniquely fixes the remaining coefficient $A = 2\eta$. This completes the proof of the uniqueness map (11).

4.1.3. Step 3: Establish the Invariant Relation and Positivity

With the dynamical results $S^\mu = 0$ and $q_{\lambda\mu\nu} = 0$, the standard quadratic identity for the torsion scalar (from Appendix B) simplifies dramatically: $T^A{}_{BC} T_A{}^{BC} = \frac{2}{3} T_\mu T^\mu$. Applying the PT projector and using the on-shell trace-lock result $T_\mu = 3\eta \partial_\mu \epsilon$, we find

$$\Pi_{\text{PT}}[T^A{}_{BC} T_A{}^{BC}] = \frac{2}{3} \Pi_{\text{PT}}[(3\eta \partial_\mu \epsilon)(3\eta \partial^\mu \epsilon)] = 6\eta^2 \Sigma_\epsilon. \quad (19)$$

Using our normalized definition of the invariant $I_T \equiv -\Pi_{\text{PT}}[T^A{}_{BC} T_A{}^{BC}]$ (see Section 2.6), this immediately yields the algebraic relation $I_T = -6\eta^2 \Sigma_\epsilon$, proving (12). The physical requirement that the kinetic term for tensor modes is positive (established in Section 6) fixes the sign choice $\eta > 0$. This concludes the proof of Theorem 2.

4.2. Implications and Diagnostics

The uniqueness theorem C1 is a cornerstone of our framework. It demonstrates that the axial and tensor components of torsion are not merely suppressed, but are algebraically eliminated from the observable sector by the dynamics. This is a powerful mechanism for avoiding the extra propagating degrees of freedom that often plague theories with torsion.

Corollary (Basis Reduction).

As a direct consequence, the "pre-lock" operator basis $\{I_1, I_2, I_3\}$ from Section 3.2 dynamically collapses into a "post-lock" basis where the physical sector is one-dimensional. The relation (12) shows that the physics can be described by the single independent invariant Σ_ϵ (or equivalently I_T)

at this order. This profound simplification is the key to the bulk equivalences we will explore in the next section.

Observational Diagnostic.

This theorem offers a sharp, falsifiable prediction. Any future detection of physical effects stemming from axial or tensor components of torsion would invalidate our framework at this order. Furthermore, the perfect alignment of T_μ and $\partial_\mu \epsilon$ can be tested in numerical simulations. The diagnostic plot in Figure 4 shows the distribution of the alignment angle, which is predicted to be sharply peaked at zero on admissible domains.




Figure 4. Alignment-angle diagnostic. Distribution of the alignment angle θ between the torsion trace T_μ and its source gradient $\partial_\mu \epsilon$, defined by $\cos \theta \equiv T_\mu \partial^\mu \epsilon / \sqrt{(T_\nu T^\nu)((\partial_\rho \epsilon)(\partial^\rho \epsilon))}$. The uniqueness map (11) predicts $\theta \simeq 0$ up to finite-domain or boundary remainders. [nb: fig_c1_alignment.py]

5. The Second Pillar: Equivalence of Constructions (C2)

Positioning & Scope. We establish bulk equivalence (modulo improvements) among three constructions. The equivalence is operational at quadratic order and on admissible domains; it is not a claim of all-order, off-shell identity. The flux-ratio diagnostic makes boundary conventions explicit, supporting reproducibility while keeping the focus on bulk dynamics.

Having established in Section 4 that our symmetric posture uniquely fixes torsion to a pure-trace form (Result C1), we now reveal a key structural property of the framework. This section presents our second main result (C2): a remarkable equivalence between three distinct methods for constructing a torsion-based gravitational action. We will demonstrate that a rank-one determinant (ROD) route, a closed-metric deformation, and a PT-projected Chern-Simons term all collapse to the *exact same bulk dynamics* at quadratic order. This non-trivial equivalence is a direct consequence of the underlying symmetries and the C1 result. It demonstrates the internal coherence of the theory and, as shown in Section 6, provides the essential ingredients needed to achieve exact luminality without fine-tuning.

Scope and Conventions.

Throughout this section, "equivalence" refers to the *equality of the bulk action at quadratic order, modulo boundary/improvement terms*. All statements are made within the A1–A6 posture and after the enforcement of the pure-trace condition C1. For mathematical convenience, we adopt the sign-compensated notation $\text{sgn}(\Sigma_\epsilon) I_T \equiv \sigma_\epsilon I_T = -6\eta^2 |\Sigma_\epsilon|$.

5.1. Three Equivalent Routes to the Same Bulk Action

We now demonstrate the equivalence by analyzing three methods for constructing a PT-even, projective-invariant theory at the quadratic order. The equivalence hinges on the key algebraic relation derived from C1, which connects the torsion scale τ to the invariant $\text{sgn}(\Sigma_\epsilon) I_T$:

$$\tau^2 = 9\eta^2 |\Sigma_\epsilon| = \frac{3}{2} (-\text{sgn}(\Sigma_\epsilon) I_T). \quad (20)$$

This identity is the algebraic engine that drives the three constructions to the same physical result.

Route 1: The Rank-One Determinant (ROD) Route

Inspired by DBI-type actions, this route constructs a Lagrangian from the determinant of a rank-one deformation of the metric involving the canonical traceless tensor $\mathcal{T}^\mu{}_\nu$ defined in Section 3. The Lagrangian is given by

$$\mathcal{L}_{\text{ROD}} = \sqrt{-g} \left(\sqrt{\det[1 + \frac{2}{3} \lambda \tau \mathcal{T}]} - 1 \right). \quad (21)$$

This construction is not Born-Infeld gravity, but shares its geometric spirit.

Route 2: The Closed-Metric (CM) Route

This route defines a new effective metric $\tilde{g}_{\mu\nu}$ via a rank-one shift, $\tilde{g}_{\mu\nu} = g_{\mu\nu} + (\frac{2}{3})\lambda \tau \hat{n}_\mu \hat{n}_\nu$. The Lagrangian is then constructed from the difference in the volume elements, $\mathcal{L}_{\text{CM}} \equiv \sqrt{-\tilde{g}} - \sqrt{-g}$.

Route 3: The PT-Even Chern-Simons (CS⁺) Route

This route starts with the topological Nieh-Yan 4-form, $\text{NY} \equiv d(e^A \wedge T_A)$. Using the identity $T^A \wedge T_A = \text{NY} + e^A \wedge e^B \wedge R_{AB}$, we apply the Hodge star and our scalar PT projector. The term from NY becomes a boundary term (by A5), and the remaining PT-even bulk piece constitutes the third route, $\mathcal{L}_{\text{CS}^+}$.

The following proposition formalizes the result that all three distinct routes lead to the same bulk action.

Proposition 1 (Three-Route Bulk Equivalence). *Under the A1–A6 posture and after enforcing C1, the quadratic-order bulk actions derived from the ROD, CM, and CS⁺ routes are identical. Specifically, each route yields a bulk Lagrangian of the form*

$$\delta^2 \mathcal{L}_X = A_\star \sqrt{-g} \text{sgn}(\Sigma_\epsilon) I_T \pmod{\nabla_\mu J_X^\mu}, \quad \text{with a universal coefficient } A_\star = \lambda^2/8, \quad (22)$$

where $X \in \{\text{ROD}, \text{CM}, \text{CS}^+\}$.

Sketch of Proof. The equivalence of the ROD and CM routes follows from their shared Jacobian, $\sqrt{-\tilde{g}}/\sqrt{-g} = \sqrt{\det[1 + (\frac{2}{3})\lambda\tau\mathcal{T}]}$. Expanding the determinant to second order using the tracelessness of $\mathcal{T}$ yields $\delta^2 \mathcal{L} \propto -\lambda^2 \tau^2$. Applying the key relation (20) then gives the result. The CS⁺ route involves a more detailed calculation with the Hodge star, but after projection and application of C1, it reduces to the same bulk term (details in Appendix C). $\square$

This remarkable equivalence, numerically verified in Figure 5, establishes the main bulk identity of this section:

$$\delta^2 \mathcal{L}_{\text{ROD}} \equiv \delta^2 \mathcal{L}_{\text{CM}} \equiv \delta^2 \mathcal{L}_{\text{CS}^+} \equiv (\lambda^2/8) \sqrt{-g} \text{sgn}(\Sigma_\epsilon) I_T \pmod{\nabla_\mu J^\mu}. \quad (23)$$




Figure 5. Residual scan for the three-route equivalence. Quadratic reductions of the ROD, CM, and CS^+ routes are compared against the target bulk action $\delta^2 \mathcal{L}_{\text{target}} = (\lambda^2/8)\sqrt{-g} \text{sgn}(\Sigma_\epsilon) I_T$. The vertical axis shows the residual after subtraction. All three routes saturate the target within numerical tolerance, confirming the bulk equivalence. *Analysis is at quadratic-order and bulk-only (mod $\nabla_\mu J^\mu$; representative Lagrangians are listed in Appendix C).* [nb: fig_c2_coeff_compare.py]

5.2. Boundary Terms and the Flux-Ratio Diagnostic

The only differences between the three routes lie in their respective boundary terms, which can be expressed as different improvement currents J_X^μ . On admissible spacetimes (as per A4), these boundary terms do not affect the bulk equations of motion. A sharp, quantitative test of their physical equivalence is provided by the flux-ratio diagnostic, which compares the integrated flux of these currents over a closed boundary $\partial\mathcal{D}$. For any two routes X and Y , our posture (A1–A5) implies that this ratio must be unity:

$$\mathcal{R}_{X/Y}[\partial\mathcal{D}] \equiv \frac{\int_{\partial\mathcal{D}} J_X^\mu d\Sigma_\mu}{\int_{\partial\mathcal{D}} J_Y^\mu d\Sigma_\mu} = 1 \quad (\text{at quadratic order}). \quad (24)$$

As shown in Figure 6, numerical computations on finite FRW balls confirm that this ratio converges to unity as the domain size increases. This demonstrates that the routes are not just bulk-equivalent, but fully physically equivalent within our posture.




Figure 6. Flux-ratio diagnostic for physical equivalence. Boundary flux ratios $\mathcal{R}_{X/Y}(R)$ for representative pairs on finite FRW balls converge to 1 as the radius R grows, in agreement with Equation (24). This confirms that the different boundary terms are physically equivalent on admissible domains. Error bars reflect the numerical tolerance propagated to the boundary terms. *Analysis is at quadratic-order (mod $\nabla_\mu J^\mu$; improvement currents listed in Appendix C).* [nb: fig_flux_ratio.py]

5.3. Implications of the Equivalence

The three-route equivalence (C2) serves as a powerful consistency check, demonstrating the robustness and internal coherence of the framework. More importantly, it furnishes us with a set of distinct but bulk-equivalent building blocks ($\mathcal{L}_{\text{ROD}}, \mathcal{L}_{\text{CM}}$). While identical in the bulk, their full expressions (including boundary terms) couple differently to metric perturbations. This distinction is the crucial ingredient that enables the coefficient-locking mechanism for achieving exact luminality, as we will demonstrate in the next section.

6. The Final Pillar: Coefficient Locking and Exact Luminality (C3)

Positioning & Scope. *The locking condition eliminates spurious TT–nonTT mixing and fixes a unique linear combination of bulk-equivalent routes. Under the admissible-domain posture, this yields the equal-coefficient identity $K = G$ and thus $c_T = 1$ at quadratic order. The point is structural sufficiency—no parameter tuning—within the stated assumptions; beyond-quadratic or non-admissible cases are outside our present scope.*

This section presents the third and central result of our framework (C3). We leverage the bulk equivalence established in Section 5 to achieve our primary goal: ensuring a luminal speed for gravitational waves as a consequence of symmetry. The strategy is to construct a general action from a linear combination of two bulk-equivalent routes (ROD and CM) and then impose a single, physically motivated condition: the absence of unphysical mixing between propagating tensor modes and non-dynamical fields. We show that this condition leads to a novel *coefficient locking* mechanism, which uniquely fixes the relative weighting of the two routes. This "locked" theory is structurally special: it satisfies an identity that forces the kinetic and gradient terms of the tensor action to be equal, thereby enforcing $c_T = 1$ by construction.

6.1. The Locking Posture: A Linear Combination of Equivalents

Based on the bulk equivalence (C2), we consider a general action constructed as a linear combination of the ROD and CM Lagrangians:

$$\mathcal{L}_{\text{chain}}(w) \equiv w_{\text{ROD}} \mathcal{L}_{\text{ROD}} + w_{\text{CM}} \mathcal{L}_{\text{CM}}, \quad w_{\text{ROD}}, w_{\text{CM}} \in \mathbb{R}. \quad (25)$$

Although these terms are identical in the bulk, their full expressions, including boundary terms, couple differently to metric perturbations. Our task is to find the specific ratio $w_{\text{ROD}} : w_{\text{CM}}$ that yields a physically consistent theory.

6.2. The Locking Condition and its Solution

We expand the total action, $\mathcal{L}_{\text{EH}} + \mathcal{L}_{\text{chain}}(w)$, to second order in ADM perturbations. The resulting Lagrangian for the tensor sector schematically takes the form

$$\delta^2 \mathcal{L} \propto \left[K(w) h_{\text{TT}}^2 - G(w) (\nabla h_{\text{TT}})^2 \right] + (h_{\text{TT}} \cdot \mathbf{M}(w) \cdot \Phi) + \dots, \quad (26)$$

where the first bracket contains the kinetic (K) and gradient (G) terms for the tensor modes h_{TT} , defining the tensor speed $c_T^2 = G(w)/K(w)$. The second term represents a potential mixing between h_{TT} and non-propagating fields Φ (such as perturbations of the lapse and shift).

Such mixing is unphysical in the sense that it implies the metric perturbation $h_{\mu\nu}$ is not a pure spin-2 eigenstate, thereby complicating its coupling to matter. To ensure that the metric fundamentally represents pure tensor propagation, we demand **no kinetic mixing**:

$$\mathbf{M}(w) = 0. \quad (27)$$

This condition translates into a 2×2 linear system for the weights $(w_{\text{ROD}}, w_{\text{CM}})$:

$$\begin{pmatrix} \mu_{\text{ROD}}^{(N)} & \mu_{\text{CM}}^{(N)} \\ \mu_{\text{ROD}}^{(\nabla N)} & \mu_{\text{CM}}^{(\nabla N)} \end{pmatrix} \begin{pmatrix} w_{\text{ROD}} \\ w_{\text{CM}} \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}. \quad (28)$$

As shown in Appendix D, the determinant of the coefficient matrix is non-zero on any generic, non-trivial background ($\Delta \neq 0$). This guarantees that the linear system has a unique, one-dimensional solution space, which we call the "locked" solution. The non-trivial solution fixes a unique ratio for the weights:

$$\boxed{\frac{w_{\text{ROD}}^*}{w_{\text{CM}}^*} = -\frac{\mu_{\text{CM}}^{(N)}}{\mu_{\text{ROD}}^{(N)}} = -\frac{\mu_{\text{CM}}^{(\nabla N)}}{\mu_{\text{ROD}}^{(\nabla N)}}.} \quad (29)$$

6.3. Emergence of Luminality from the Locked Solution

This specific, "locked" combination of weights, denoted by w^* , is structurally remarkable. A direct calculation (detailed in Appendix D) reveals that for this precise choice, the kinetic and gradient coefficients satisfy a powerful *equal-coefficient identity*:

$$\boxed{K(w^*) - G(w^*) = a^{-3} \partial_\mu (a^3 J_\Delta^\mu).} \quad (30)$$

The right-hand side is a total divergence, which, under the admissible boundary conditions of our posture (A4), does not contribute to the equations of motion. This identity dynamically enforces the equality of the kinetic and gradient kernels:

$$K(w^*) = G(w^*) \implies \boxed{c_T^2 = \frac{G(w^*)}{K(w^*)} = 1.} \quad (31)$$

This proves our main theorem for this section:

Theorem 3 (Coefficient Locking Ensures Exact Luminality (C3)). *Within the A1–A6 posture, demanding the absence of TT–nonTT mixing uniquely fixes the theory to a "locked" state w^* . In this state, the equal-coefficient identity (30) holds, which structurally ensures that the tensor speed is exactly luminal ($c_T = 1$) at*

quadratic order. The resulting tensor action is identical to that of General Relativity, propagating only two degrees of freedom.

The consequences are visualized in Figure 7, which shows the tensor speed deviation across the parameter space of weights; the speed becomes exactly luminal only along the "locking curve" defined by our condition. Figure 8 further contrasts the strictly luminal propagation in the locked theory with the non-luminal behavior of unlocked choices. In summary, our framework does not just permit $c_T = 1$, it enforces it through a structural mechanism born from symmetry and consistency.




Figure 7. Heatmap of the tensor-speed deviation. Representative scan of $|c_T - 1|$ over the $(w_{\text{ROD}}, w_{\text{CM}})$ plane on an admissible background (A1–A5, C1 in force). The *locking curve* defined by (28) is overlaid. Along this curve, the unphysical TT–nonTT mixing vanishes and, as a consequence, $c_T = 1$ holds (31). [nb: fig_c3_cT_heatmap.py]




Figure 8. Tensor dispersion $c_T(k)$: locked vs. unlocked. Comparison of $c_T(k)$ for the *locked* weight ratio w^* from (29) (solid line, showing $c_T = 1$) and representative *unlocked* choices (dashed line). [nb: fig_c3_dispersion.py]

7. The Locked Theory: Action, Luminality, and Degrees of Freedom

Positioning & Scope. We show that, at this order and within the posture, the tensor sector matches GR and propagates two degrees of freedom. The Hamiltonian analysis is framed to make the constraint algebra and degree-of-freedom counting explicit. Extensions to matter-rich or higher-order settings are deferred to future work.

Having established the final link in our logical chain—the coefficient locking mechanism (C3)—we now consolidate the physical properties of the resulting "locked" theory. This section presents the final form of the tensor action, confirms its properties, and verifies that it describes the correct number of physical degrees of freedom. We first re-state the central equal-coefficient identity that enforces luminality. We then perform a Hamiltonian analysis to confirm that the theory is well-behaved, free of ghosts, and propagates precisely two tensor modes, with no additional unphysical states.

Throughout this section, we work within the full A1–A6 posture, with the C1 (pure-trace), C2 (equivalence), and C3 (locking) results all in force.

7.1. The Equal-Coefficient Identity and the Locked Tensor Action

Let us consider the total action $\mathcal{L}_{\text{tot}} = \mathcal{L}_{\text{EH}}(e, \omega) + \mathcal{L}_{\text{chain}}(w^*)$, expanded to quadratic order in perturbations. As shown in Section 6, the coefficient locking fixes the weights to w^* , which eliminates the mixing between tensor (TT) and non-tensor modes. The action for the tensor sector is

$$\delta^2 \mathcal{L}_{\text{TT}} = \frac{M_{\text{Pl}}^2}{8} a^3 \left[K(w^*) h_{ij}^{\text{TT}} h_{ij}^{\text{TT}} - G(w^*) \frac{(\partial_k h_{ij}^{\text{TT}})^2}{a^2} \right]. \quad (32)$$

The key result, which is a direct consequence of the C1 and C2 results, is the following *total-divergence identity*:

$$K(w^*) - G(w^*) = a^{-3} \partial_\mu (a^3 J_\Delta^\mu). \quad (33)$$

Here, J_{Δ}^{μ} is a quadratic improvement current (see Appendix D). Since the right-hand side is a total divergence, it does not contribute to the dynamics on an admissible domain (A4). This identity therefore dynamically enforces the equal-coefficient condition:

$$K(w^*) = G(w^*) \quad \implies \quad c_T^2 = \frac{G(w^*)}{K(w^*)} = 1 \text{ (exact at quadratic order)}. \quad (34)$$

This is a central result of our framework: exact luminality is not an assumption or a fine-tuning, but a structural consequence of the theory's symmetries. Imposing the standard GR normalization for the tensor action then yields the final, locked action:

$$\delta^2 S_{\text{TT}} = \frac{M_{\text{Pl}}^2}{8} \int d^4x a^3 \left[\dot{h}_{ij}^{\text{TT}} \dot{h}_{ij}^{\text{TT}} - \frac{(\partial_k h_{ij}^{\text{TT}})^2}{a^2} \right], \quad c_T^2 = 1. \quad (35)$$

At this order, the resulting tensor dynamics are indistinguishable from those of General Relativity, as confirmed by the waveform comparison in Figure 9.




Figure 9. GW waveform overlay (GR vs. locked). Time-domain comparison of a representative TT mode in GR (reference) and in the *locked* theory (this work). The two traces are indistinguishable within numerical tolerance, consistent with the fact that the locked action (35) is identical to that of GR at this order. [nb: fig_gw_waveform_overlay.py]

7.2. Hamiltonian Analysis and Degrees of Freedom

To complete our analysis, we verify that the theory is dynamically consistent and propagates the correct number of degrees of freedom (DoF). We perform a Hamiltonian analysis using the standard ADM (3+1) decomposition.

Constraint Analysis.

The full set of configuration variables includes the metric components (h_{ij}, N, N^i) and the algebraic fields associated with torsion. The analysis of the constraint structure reveals:

- **Primary Constraints:** The lapse N and shift N^i are auxiliary fields, leading to the primary constraints $\pi_N \approx 0$ and $\pi_i \approx 0$. The Lagrange multiplier Λ^μ that enforces the trace lock is also non-dynamical, yielding $p_{\Lambda^\mu} \approx 0$.
- **Secondary Constraints:** The algebraic nature of the Palatini equations of motion (from varying ω) and the trace-lock constraint (from varying Λ^μ) ensures that no new propagating modes are introduced. These constraints algebraically eliminate all non-GR torsion components ($S^\mu =$

$0, q_{\lambda\mu\nu} = 0$) and fix the trace component ($T_\mu = 3\eta\partial_\mu\epsilon$). The preservation of primary constraints in time generates the standard Hamiltonian and momentum constraints of GR, $\mathcal{H} \approx 0$ and $\mathcal{H}_i \approx 0$.

Constraint Algebra and DoF Count.

Crucially, all constraints related to the torsion sector are algebraic and can be solved explicitly, leaving no residual dynamics. On an admissible domain (A4), the remaining constraints ($\pi_N, \pi_i, \mathcal{H}, \mathcal{H}_i$) form a first-class algebra identical to that of GR at this order. A standard DoF count for the metric sector confirms that the theory propagates the correct number of physical modes:

$$N_{\text{DoF}} = \frac{1}{2}((\text{phase space variables}) - 2 \times (\text{first-class constraints})) = \frac{1}{2}((6 + 6) - 2 \times 4) = 2.$$

The theory propagates precisely two degrees of freedom, corresponding to the two tensor polarizations of gravitational waves. This is confirmed numerically by diagonalizing the kinetic kernel and counting the number of non-zero eigenvalues, as shown in Figure 10.




Figure 10. DoF spectrum (eigenvalue stem plot). Eigenvalues of the quadratic kinetic kernel after integrating out non-propagating variables. The count of eigenvalues above the numerical tolerance (`deg_tol`) tracks the number of physical degrees of freedom, $\#phys \approx 2$, across a range of background configurations. This confirms the absence of extra propagating modes at quadratic order. [nb: `fig_c3_degeneracy.py`]

7.3. Summary of the Locked Theory

This section has consolidated the main results of our framework. The coefficient locking mechanism leads to an equal-coefficient identity, which enforces exact luminality ($c_T = 1$) for tensor modes as a structural consequence of the underlying symmetries. The resulting effective action for the tensor sector is shown to be identical to that of General Relativity at this order. A Hamiltonian analysis confirms that the theory is well-posed, with a first-class constraint algebra that removes all non-physical modes, leaving precisely two propagating tensor degrees of freedom. The framework thus provides a consistent, predictive, and structurally luminal theory of gravity.

8. Phenomenological Consequences and Observational Tests

Positioning & Scope. *Our data-facing statements are deliberately narrow: (i) minimal coupling to Dirac fermions yields a clean tensor sector at this order; (ii) the unique NLO dispersion $\propto k^2$ defines a distinctive,*

falsifiable target across PTA/LISA/LVK bands; (iii) we propose practical reporting via b/Λ^2 . A full cosmological parameter analysis is outside the present scope.

Having established the core theoretical properties of the framework—unique torsion (C1), bulk equivalence (C2), and locked luminosity (C3)—we now turn to its phenomenological consequences. An internally consistent gravitational theory must also interact cleanly with matter and produce testable predictions. This section addresses these points. We first examine the coupling to Dirac fermions, showing that it is trivial at leading order. We then derive the unique next-to-leading-order (NLO) correction to the tensor dispersion relation, which stands as the primary observational signature of this work. Finally, we outline a clear, data-driven strategy to test, constrain, or falsify the theory.

Our entire phenomenological analysis is built upon the cornerstone result C1 (Theorem 2), which dictates that the only surviving component of torsion in the observable sector is the pure-trace part aligned with the Stueckelberg gradient:

$$T_\mu = 3\eta \partial_\mu \epsilon, \quad S^\mu = 0, \quad q_{\lambda\mu\nu} = 0. \quad (36)$$

All simplifications that follow are direct consequences of this dynamical result, not additional assumptions.

8.1. Interaction with Dirac Fermions: A Null Result at Leading Order

We begin by considering a standard Dirac spinor minimally coupled to the Riemann–Cartan geometry. The interaction Lagrangian between the spinor’s vector current (J^μ) and axial-vector current (J_5^μ) and the torsion irreps is given by (see Appendix E for conventions):

$$\mathcal{L}_{\psi,T} = e \left(c_A S_\mu J_5^\mu + c_V T_\mu J^\mu \right) + \text{boundary terms}. \quad (37)$$

Our framework makes two powerful simplifying predictions regarding this interaction.

No Axial Channel.

The uniqueness theorem C1 directly enforces $S^\mu = 0$. Consequently, the axial coupling term, which is often tightly constrained by experiment, vanishes identically at tree level: $e c_A S_\mu J_5^\mu \equiv 0$.

Removable Trace Channel.

The only remaining interaction is the trace channel, which by C1 takes the form $\mathcal{L}_{\psi,T} \propto e (\partial_\mu \epsilon) J^\mu$. This coupling, however, is not a fundamental interaction but an artifact of the chosen field basis. As detailed in Appendix E, it can be completely removed by a local, anomaly-free vector phase redefinition of the Dirac field ($\psi \mapsto e^{i\alpha\epsilon} \psi$), which reduces the term to a harmless boundary improvement.

In summary, the symmetric posture of our theory ensures a minimally coupled fermion sector at this order: both potential torsion-fermion interaction channels are either absent by dynamics or removable by a field redefinition.

8.2. The Unique NLO Signature: A k^2 Tensor Dispersion

Beyond its leading-order luminosity, our framework makes a sharp and unique prediction at the next-to-leading order (NLO) in the derivative expansion. The same symmetries that enforce $c_T = 1$ at leading order permit only one new bulk operator that can affect the tensor dispersion.³ This operator,

³ Other admissible NLO scalars either renormalize the kinetic and gradient terms K and G equally (leaving c_T unchanged), or can be absorbed into boundary improvements under assumption A4.

which can be understood from an effective field theory (EFT) perspective (see Appendix G), leads to the dispersion relation:

$$\omega^2 = k_{\text{phys}}^2 \left[1 + b \frac{k_{\text{phys}}^2}{\Lambda^2} \right] \implies \boxed{c_T^2(k) \equiv \frac{\omega^2}{k_{\text{phys}}^2} = 1 + b \frac{k_{\text{phys}}^2}{\Lambda^2}} \quad (k_{\text{phys}} \ll \Lambda), \quad (38)$$

where b is a dimensionless coefficient of $\mathcal{O}(1)$ and Λ is the effective scale of new physics. This quadratic dependence on frequency is the defining, distinctive signature of our framework at NLO, as visualized in the log-log plot of Figure 11.




Figure 11. NLO signature and slope fit. The predicted tensor-speed deviation $\delta c_T^2(k)$ from the locked leading-order value $c_T^2 = 1$, plotted as a function of physical wavenumber k . The log-log plot shows the characteristic slope $\hat{n} \simeq 2$, a direct consequence of the unique NLO operator. The vertical offset of the line fixes the ratio b/Λ^2 . [nb: fig_nlo_offsets.py]

8.3. Data-Facing Strategy: Constraints and Falsification

The NLO prediction (38) provides a direct and powerful bridge to observational data.

Recipe for Constraining the Theory.

Any observational constraint on c_T can be translated into a bound on the fundamental parameter space of the theory. Given a measurement or an upper limit on $\delta c_T^2(f_*)$ in a frequency band centered at f_* , we propose reporting a standardized constraint on the combination b/Λ^2 :

$$\frac{b}{\Lambda^2} = \frac{\delta c_T^2(f_*)}{(2\pi f_*)^2}. \quad (39)$$

This provides a simple and unified way to test the framework across different experiments (e.g., PTA, LISA, LVK) and to combine their results. For instance, the existing bound from GW170817 implies a scale of $\Lambda \gtrsim 9.3 \times 10^{-6}$ eV (for $|b| \sim 1$).

Null Tests for Foundational Assumptions.

A robust framework should also offer ways to test its own foundational assumptions. Our posture provides two independent null tests for the spurion limit of the field ϵ :

1. **The Slope Test:** The theory predicts a pure power-law slope of 2 for the dispersion relation in a log-log plot. A statistically significant deviation from this slope would indicate residual dynamics in the ϵ field, thus falsifying the strict spurion limit assumed in our leading-order analysis.
2. **The Equivalence Test:** The bulk equivalence of the ROD and CM routes (Result C2) is a direct consequence of the spurion limit. Any non-zero measurement of the "route-difference coefficient," $\Delta A_{\star}(k) \equiv A_{\star}^{(\text{rank-one determinant route})}(k) - A_{\star}^{(\text{CM})}(k)$, would provide direct evidence for new dynamics beyond our posture.

Passing both null tests would provide strong corroborating evidence for the framework's foundational assumptions, while failing either would point towards new, potentially observable physics. This combination of a unique observational signature and clear falsification criteria makes the framework highly testable.

9. Reproducibility Supplement

Positioning & Scope. *This lean supplement points to a public repository with figure generators, configuration files, tests, and checksums. The intent is to enable independent verification of the main claims (C1)–(C3) without embedding long code listings in the manuscript, aligning with the editorial focus of the present framework on soundness and clarity.*

This paper's core claims (C1, C2, C3) are supported by analytical and numerical calculations, which are made fully reproducible through a public, open-source repository. All code, figure generators, configuration files, and validation tests are available at:

<https://github.com/ice91/palatini-pt-spurion/>

This supplement provides a minimal, repository-backed map for readers to rebuild all figures and independently validate the paper's main results. To ensure clarity and conciseness, we avoid embedding large code blocks; the repository is self-contained and fully documented.

9.1. Repository Layout and Terminology

Directory Structure.

The repository is organized into the following top-level directories: `scripts/` (figure generators), `configs/` (numerical grids & coefficient files in JSON format), `palatini_pt/` (core library), `tests/` (pytest validation suite), `figs/` (output figures and data artifacts), and `notebooks/` (exploratory script mirrors). A one-shot driver script, `scripts/make_all_figs.py`, rebuilds all paper figures from scratch.

Terminology Note.

The paper uses the term *Rank-One Determinant (ROD) route*. For historical reasons, the corresponding configuration files in the repository use the legacy token `dbi` (e.g., `configs/coeffs/dbi.json`). They refer to the same construction.

9.2. Validation and Verification Workflow

A reader can reproduce all results and validate the main claims in three simple steps:

1. **Setup Environment (conda/mamba):** The exact computational environment is pinned in the `environment.yml` file. It can be created with a single command:

```
conda env create -f environment.yml; conda activate palpt
```
2. **Rebuild All Figures:** All figures in this paper can be regenerated by running the master script:

```
python scripts/make_all_figs.py
```
3. **Validate Claims (C1, C2, C3):** A comprehensive test suite using pytest is provided to programmatically verify the core results. This includes tests for the uniqueness of torsion (C1), the bulk equivalence of the three routes (C2), the coefficient locking (C3), and the NLO dispersion. The suite can be run with:

```
pytest -q
```

9.3. Figure and Artifact Map

Table 2 provides a map from each figure in the paper to its generating script and required configuration files.

Table 2. Figure file map (repository-backed).

ID	Generator (scripts/)	Inputs (configs/)	Output (PDF under figs/pdf/)
Figure 2	fig_c1_pure_trace.py	—	figs/pdf/fig1_c1_pure_trace.pdf
Figure 4	fig_c1_alignment.py	—	figs/pdf/fig2_c1_alignment.pdf
Figure 5	fig_c2_coeff_compare.py	coeffs/dbi.json, coeffs/closed.json, coeffs/cspp.json	figs/pdf/fig3_c2_coeff_compare.pdf
Figure 7	fig_c3_cT_heatmap.py	optional: paper_grids.yaml	figs/pdf/fig4_c3_cT_heatmap.pdf
Figure 8	fig_c3_dispersion.py	optional: paper_grids.yaml	figs/pdf/fig5_c3_dispersion.pdf
Figure 10	fig_c3_degeneracy.py	—	figs/pdf/fig6_c3_degeneracy.pdf
Figure 9	fig_gw_waveform_overlay.py	—	figs/pdf/fig7_gw_waveform_overlay.pdf
Figure 11	fig_nlo_offsets.py	optional: paper_grids.yaml	figs/pdf/fig8_nlo_offsets.pdf
Figure 6	fig_flux_ratio.py	optional: paper_grids.yaml	figs/pdf/fig9_flux_ratio.pdf

9.4. Data Integrity and Version Control
Checksums.

Every generated PDF and data artifact is shipped with a corresponding .md5 sidecar file (e.g., figs/pdf/fig1_c1_pure_trace.pdf.md5). Integrity can be verified using a standard checksum utility: `md5sum -c figs/pdf/*.md5`

Version Pinning.

To ensure long-term reproducibility, we cite the exact Git revision hash used to generate the final paper artifacts and tag the corresponding release. The repository provides snapshots of the figures (figs.tar.gz) that match the committed artifacts. The results are deterministic on the platforms listed in the repository’s README.md file, requiring no external accelerators or downloads.

Summary. This approach to reproducibility, centered on a public, pinned repository with scripted generation, configuration-controlled coefficients, a comprehensive test suite, and verifiable checksums, ensures that the paper’s claims can be scrutinized and validated by the community with minimal friction.

10. Context and Relation to Other Frameworks

Positioning & Scope. *We position the posture relative to EC/MAG traditions and Palatini-type modifications, highlighting how the scalar PT projection and projective invariance address known pitfalls at quadratic order. This is not an exhaustive survey; references are curated to clarify the structural differences most relevant to our claims.*

This section places our *scalar-PT projected Palatini posture* in the context of several related research areas: (i) the historical Einstein-Cartan/metric-affine (EC/MAG) tradition, (ii) Palatini-type modified gravity and its known pitfalls, and (iii) the post-GW170817 observational landscape. We conclude with a set of compact, operational notes that allow the main claims (C1–C3) to be checked independently, emphasizing that all results are restricted to the posture defined by A1–A6.

10.1. Historical and Geometric Context (EC/MAG; Metric-Affine)

The decomposition of torsion into its irreducible trace, axial, and traceless components, and the independent variation of the vierbein (e) and spin connection (ω), are standard practices in the EC/MAG tradition [6,8]. Our use of the Nieh–Yan 4-form for parity bookkeeping is also conventional [9–12]. Boundary and improvement terms are handled within established covariant phase-space frameworks [14,15].

Against this backdrop, our framework introduces two key distinguishing features: we *project observables to scalar, PT-even densities* and *enforce projective invariance* via a non-dynamical Stueckelberg compensator ϵ that enters only through $\mathcal{T}_\mu \equiv T_\mu - \partial_\mu \epsilon$. Within our two-derivative posture, this approach leads directly to our main results: (C1) the dynamical elimination of axial and traceless torsion; (C2) the bulk equivalence of three distinct constructions; and (C3) the coefficient-locking mechanism that ensures luminal tensor propagation without fine-tuning.

10.2. Palatini-Type Modified Gravity and Known Challenges

Palatini-type modifications, most notably Palatini $f(R)$ gravity, have a rich history but also face well-documented challenges, particularly when matter is included. These include equivalence to constrained scalar–tensor theories, tight post-Newtonian bounds, and potential pathologies in stellar contexts [19–21]. These issues motivate our symmetry-selected approach, where the observable sector is defined *before* variation and boundary effects are explicitly accounted for.

How our posture differs.

Our framework is structurally designed to sidestep these common pitfalls at the quadratic order analyzed:

1. **Observable Projection.** The scalar-PT projector removes PT-odd pseudoscalars *before* variation, preventing contamination of the tensor sector by parity-odd densities.
2. **Projective Invariance with a Spurion Limit.** By allowing only the invariant combination $\mathcal{T}_\mu$ to enter observables, the axial and traceless torsion modes are dynamically removed (C1), thus avoiding the introduction of extra propagating degrees of freedom that are often problematic.
3. **Explicit Boundary Accounting.** Improvement currents are treated as boundary terms under A4–A5. This posture is what enables the bulk equivalence (C2) and sets the stage for the equal-coefficient identity $K - G = \partial_\mu(\dots)$ that yields $c_T = 1$ (C3).

For contrast, Chern–Simons modified gravity introduces a dynamical pseudoscalar and parity-odd effects [25,26]. Our framework, instead, operates within a parity-even, projected scalar sector with controlled boundary terms.

10.3. Post-GW170817 Constraints and Torsionful Gravitational Waves

The multimessenger observation GW170817/GRB170817A provided an extremely tight constraint on the tensor speed, strongly disfavoring theories where $c_T \neq c$ [27,31]. While much of the subsequent research focused on Horndeski/EFT models, theories with torsion have also been examined. In frameworks like Poincaré gauge gravity and Einstein–Cartan theory, tensor waves are often found to be luminal, but their amplitudes, attenuation, or polarization content can differ from GR [37,38].

Our Positioning.

Our framework provides a clean, even-parity route to exact luminality at quadratic order via the structural C3 identity, not through parameter tuning.⁴ Furthermore, it predicts a unique next-to-leading order deviation:

$$\delta c_T^2(k) = b \frac{k^2}{\Lambda^2} \quad (k \ll \Lambda),$$

which is expressly designed for multi-band tests (PTA/LISA/LVK) using the log-slope diagnostic described in Section 8. This offers a falsifiable bridge between symmetry and data: a confirmed k^2 -type dispersion would constrain (b, Λ) , while its absence in the regime where the EFT is valid would disfavor our posture.

10.4. A Checkable Summary of Core Claims

To facilitate independent verification, we provide a compact, checkable summary of our main claims and their scope.

1. **Scope.** All quadratic bulk equalities (C2) and the luminality identity (C3) are asserted *within* the A1–A6 posture. Cases with topological torsion defects or boundary conditions that inject new canonical pairs fall outside this scope.
2. **C1 (Pure-trace alignment).** In the scalar- PT projected sector at quadratic order, axial and traceless torsion vanish algebraically via the equations of motion. Only the trace component aligned with $\partial\epsilon$ remains. A robust detection of axial-torsion effects in observables would falsify C1.
3. **C2 (Three-route bulk equivalence).** The ROD, CM, and PT-even CS/Nieh–Yan routes share the same bulk coefficient $A_\star$ and differ only by improvement terms. The flux ratio diagnostic $\mathcal{R}_{X/Y}[\partial\mathcal{D}] \rightarrow 1$ on admissible domains provides a practical check.
4. **C3 (Equal-coefficient locking).** The requirement of no TT-nonTT mixing fixes the weights $w^\star$, leading to the identity $K(w^\star) - G(w^\star) = \partial_\mu(\dots)$. This enforces $c_T = 1$ at quadratic order and yields the EFT-consistent NLO dispersion.

Navigation.

Table 3 summarizes how our scalar- PT Palatini posture differs from other frameworks, while Table 4 maps our assumptions to standard concepts in the literature. The claims C1–C3 are proven in the main text under these assumptions. Observational guidance for the NLO dispersion is given in Section 8.

Table 3. At-a-glance comparison of relevant frameworks. "Bulk eq." refers to quadratic bulk action equality modulo improvements; "Obs. torsion" is the torsion content that survives in the observable tensor sector.

Framework	c_T	DoF	Obs. torsion	Parity	Notes / Refs
This work (scalar- PT Palatini)	$= 1$	2	pure trace aligned with $\partial\epsilon$ (C1)	even	bulk eq. (C2); k^2/Λ^2 NLO; C1–C3; Section 8
EC/MAG (Einstein–Cartan)	≈ 1	2 (+)	axial couples to spin; no speed shift (minimal EC)	mostly even	[6,8]
Palatini $f(R)$	≈ 1	2+scal.	typically none (torsionless)	even	[19–21]
CS modified gravity	≈ 1	2	none	odd	[25,26]
PGG	luminal?	2+	model-dep.	mixed	[37]

⁴ While the outcome $c_T = 1$ resembles that of DHOST Class I theories (where kinetic terms are degenerate), our mechanism is distinct: it arises from the PT -symmetric alignment of torsion (C1) and the algebraic locking of bulk-equivalent routes (C3), rather than from degeneracy conditions in a scalar-tensor metric sector.

Table 4. Operational mapping of the global assumptions (A1–A6) to mainstream constructs and indicative references.

Assumption	Mainstream analogue / operational meaning	Indicative refs
A1 (domain/measure)	Oriented Lorentzian patches; AF/FRW fall-offs; projector is self-adjoint on real scalar densities preserved under PT .	[14,15]
A2 ($[\mathcal{PT}, *] = 0$)	Orientation-preserving PT ; standard parity book-keeping with the Holst/Nieh–Yan split.	[9–12]
A3 (project–vary commute)	Linear, self-adjoint scalar- PT projector commutes with Palatini variation on (e, ω) when the spurion is non-dynamical.	[6,14]
A4 (boundary posture)	Improvement currents contribute no canonical pairs; flux is well-defined and slice/gauge independent at quadratic order.	[14,15]
A5 (Nieh–Yan as boundary)	$NY = d(e \wedge T)$ is treated as a boundary counterterm on admissible patches; topological caveats noted.	[9,12]
A6 (two-deriv. trunc. & spurion)	Analysis is restricted to quadratic order in fields with at most one derivative per building block. The compensator ϵ is treated as a non-dynamical background field.	(This paper; see Section 2)

Acknowledgments

The author is grateful to the anonymous referees for comments that improved the manuscript. Limited use of generative language tools was made for stylistic refinement; all scientific reasoning, derivations, and conclusions remain solely the responsibility of the author.

Statements and Declarations

Competing interests.

The author has no relevant financial or non-financial interests to disclose.

Data availability.

All code and figure-generation scripts are openly available at <https://github.com/ice91/palatini-pt-spurion/> under a permissive license; version pins and reproduction instructions are provided in Supplement R.

Code availability.

Same as Data availability.

Appendix A. Projection–Variation Commutation (A3) and Variational Identities

This appendix establishes Assumption A3 in full generality for scalar densities and compiles the variational identities for $\sqrt{-g}$, $e^{\mu\nu\rho\sigma}$, and the Hodge star $*$ under the Palatini posture with the scalar PT projector of Section 2. We also make explicit the boundary/topology posture (A4) used when trading improvements for boundary fluxes, and we record the PT selection rules used throughout the paper.

Appendix A.1. Setup and Conventions

We work with independent vierbein e^A_μ and a metric-compatible spin connection $\omega^{AB}_\mu = -\omega^{BA}_\mu$ (Palatini posture). The observable *scalar* densities are mapped to a real, PT -even sector by the projector

$$\Pi_{PT}[\mathcal{O}] \equiv \frac{1}{2}(\mathcal{O} + \mathcal{PT}[\mathcal{O}])\Big|_{\text{real scalar}}, \qquad \Pi^2_{PT} = \Pi_{PT}, \qquad (A1)$$

with the combined PT acting anti-linearly (complex conjugation accompanies T) and preserving the chosen orientation (A2), so that $[\mathcal{PT}, *] = 0$ on forms. In particular, PT is orientation-preserving and the Levi–Civita tensor $\epsilon^{\mu\nu\rho\sigma}$ is PT -even.⁵ The internal phase $\epsilon(x)$ is a *spurion*: it enters observables only via $\partial\epsilon$ and is *not* varied. All statements below are thus variations with respect to $(e, \omega, \text{matter})$ while keeping ϵ fixed.

Appendix A.2. Projector Properties: Idempotence, Self-Adjointness, and Selection Rules

Proposition A1 (Self-adjointness of Π_{PT} on real scalars). *Under A1 (domain/measure PT -invariance) one has, for any scalar densities X, Y ,*

$$\int \Pi_{PT}[X] \Pi_{PT}[Y] = \int \Pi_{PT}[XY]. \quad (\text{A2})$$

Proof. Expand the left-hand side using Equation (A1) and the fact that $\mathcal{PT}^2 = \text{id}$:

$$\int \frac{1}{4} (X + \mathcal{PT}[X]) (Y + \mathcal{PT}[Y]) = \frac{1}{4} \int (XY + X \mathcal{PT}[Y] + \mathcal{PT}[X] Y + \mathcal{PT}[X] \mathcal{PT}[Y]).$$

Using A1, $\int Z = \int \mathcal{PT}[Z]$, the cross terms rearrange into $\frac{1}{2} \int (XY + \mathcal{PT}[XY]) = \int \Pi_{PT}[XY]$. $\square$

Proposition A2 (Selection rules for the scalar projector). *With A1–A2 and metric compatibility, for any admissible tensors X :*

$$\Pi_{PT}[\epsilon^{\mu\nu\rho\sigma} X_{\mu\nu\rho\sigma}] = \begin{cases} \epsilon^{\mu\nu\rho\sigma} X_{\mu\nu\rho\sigma} \text{ (up to taking the real part),} & \text{if } X \text{ is } PT\text{-even,} \\ 0, & \text{if } X \text{ is } PT\text{-odd,} \end{cases} \quad (\text{A3})$$

$$\Pi_{PT}[T_\mu \partial^\mu \epsilon] \in \mathbb{R} \quad (T_\mu \partial^\mu \epsilon \text{ is } PT\text{-even; it survives pre-lock and maps to } 3\eta \Sigma_\epsilon \text{ post-lock}), \quad (\text{A4})$$

$$\Pi_{PT}[T^A_{BC} T^B_C] \in \mathbb{R}, \quad \Pi_{PT}[(\partial\epsilon)^2] \in \mathbb{R}. \quad (\text{A5})$$

Proof (sign count). Under A2 the chosen orientation is preserved and $[\mathcal{PT}, *] = 0$. Since PT is orientation-preserving, the Levi–Civita tensor is PT -even,

$$\mathcal{PT}[\epsilon^{\mu\nu\rho\sigma}] = \epsilon^{\mu\nu\rho\sigma},$$

so for any $X_{\mu\nu\rho\sigma}$ one has

$$\mathcal{PT}[\epsilon^{\mu\nu\rho\sigma} X_{\mu\nu\rho\sigma}] = \epsilon^{\mu\nu\rho\sigma} \mathcal{PT}[X_{\mu\nu\rho\sigma}].$$

Therefore the PT parity of the scalar density $\epsilon^{\mu\nu\rho\sigma} X_{\mu\nu\rho\sigma}$ is determined entirely by that of $X_{\mu\nu\rho\sigma}$, yielding (A3) after applying the projector definition $\Pi_{PT}[\mathcal{O}] = \frac{1}{2}(\mathcal{O} + \mathcal{PT}[\mathcal{O}])|_{\text{real}}$.

For the spurion gradient, the PT -transformation table of Section 2 assigns opposite P/T parities to $T_\mu \equiv T^v_{\mu\nu}$ and $\partial_\mu \epsilon$ so that their product $T_\mu \partial^\mu \epsilon$ is PT -even; the projector thus returns its real part and it survives the projection step. Quadratic contractions T^2 and $(\partial\epsilon)^2$ are manifestly PT -even (they are built from metric contractions of vectors and gradients), so again the projector reduces to taking the real part, consistent with Prop. A1. $\square$

Appendix A.3. Commutation of Projection with Variation (A3)

Theorem A1 (Projection–variation commutation (A3)). *Let $\mathcal{O}$ be any local scalar density built from $(e, \omega, \partial\epsilon, \dots)$. Then, for variations with respect to $(e, \omega, \text{matter})$ at fixed ϵ ,*

$$\delta \Pi_{PT}[\mathcal{O}] = \Pi_{PT}[\delta \mathcal{O}]. \quad (\text{A6})$$

⁵ We work with standard $P : (t, \vec{x}) \mapsto (t, -\vec{x})$, $T : (t, \vec{x}) \mapsto (-t, \vec{x})$, so that $PT : (t, \vec{x}) \mapsto (-t, -\vec{x})$ has $\det J_{PT} = +1$ and preserves orientation. The Levi–Civita symbol $\epsilon^{\mu\nu\rho\sigma}$ is constant, and $\epsilon^{\mu\nu\rho\sigma} = \epsilon^{\mu\nu\rho\sigma} / \sqrt{-g}$ is PT -even.

Proof. By definition, $\delta \Pi_{PT}[\mathcal{O}] = \frac{1}{2}(\delta \mathcal{O} + \delta \mathcal{PT}[\mathcal{O}])|_{\text{real}}$. It suffices to show $\delta \mathcal{PT}[\mathcal{O}] = \mathcal{PT}[\delta \mathcal{O}]$. The PT action on fields is an involutive automorphism on the local functional algebra, and it is anti-linear only through global complex conjugation (time reversal). For any complex functional F one has $\delta \bar{F} = \overline{\delta F}$ because the variation acts linearly on fields and does not act on the numerical i . Therefore, with Φ denoting collectively the fields that are varied and $\Phi^{\mathcal{PT}}$ their PT image,

$$\delta \mathcal{PT}[\mathcal{O}(\Phi)] = \delta \overline{\mathcal{O}(\Phi^{\mathcal{PT}})} = \overline{\delta \mathcal{O}(\Phi^{\mathcal{PT}})} = \mathcal{PT}[\delta \mathcal{O}(\Phi)],$$

where we used that $\mathcal{PT}$ does not touch the spurion (fixed) and commutes with derivatives and index operations under A2. Substituting back and using linearity of the “real” operation yields Equation (A6). $\square$

Remarks.

(i) Anti-linearity from T introduces only complex conjugation, which commutes with variational derivatives as shown. (ii) The assumption that ϵ is not varied (spurion posture) is essential; if one promotes ϵ to a dynamical field, additional boundary terms appear but Equation (A6) continues to hold for the *scalar* projector provided the same posture (A1–A2) is kept for the extended field space.

Appendix A.4. Variational Identities for $\sqrt{-g}$, $\epsilon^{\mu\nu\rho\sigma}$, and the Hodge Star

We collect formulas used repeatedly in Secs. 2–7. We write $h_{\mu\nu} \equiv \delta g_{\mu\nu}$ and $h \equiv g^{\mu\nu} h_{\mu\nu}$.

Metric and vierbein.

With $g_{\mu\nu} = \eta_{AB} e^A_\mu e^B_\nu$,

$$\delta g_{\mu\nu} = \eta_{AB} (e^A_\mu \delta e^B_\nu + e^A_\nu \delta e^B_\mu), \quad \delta g^{\mu\nu} = -g^{\mu\alpha} g^{\nu\beta} \delta g_{\alpha\beta}. \quad (\text{A7})$$

Determinant and Levi–Civita tensor.

$$\delta \sqrt{-g} = \frac{1}{2} \sqrt{-g} h, \quad \delta \epsilon_{\mu\nu\rho\sigma} = \frac{1}{2} \epsilon_{\mu\nu\rho\sigma} h, \quad \delta \epsilon^{\mu\nu\rho\sigma} = -\frac{1}{2} \epsilon^{\mu\nu\rho\sigma} h. \quad (\text{A8})$$

These follow from $\epsilon_{\mu\nu\rho\sigma} = \sqrt{-g} \epsilon_{\mu\nu\rho\sigma}$ and $\epsilon^{\mu\nu\rho\sigma} = \epsilon^{\mu\nu\rho\sigma} / \sqrt{-g}$, with $\epsilon_{\mu\nu\rho\sigma}$ the (constant) Levi–Civita symbol.

Hodge star.

Let α be a p -form and $h^\mu_\nu \equiv g^{\mu\lambda} h_{\lambda\nu}$. Then the variation of $*$ with respect to h is

$$\delta(*\alpha) = *(\delta\alpha) + \frac{1}{2} h(*\alpha) - \frac{1}{2} *(H \cdot \alpha), \quad (H \cdot \alpha)_{\mu_1 \dots \mu_p} \equiv \sum_{i=1}^p h_{\mu_i}^{\nu} \alpha_{\mu_1 \dots \nu \dots \mu_p}. \quad (\text{A9})$$

In particular, for 2-forms F (frequent in the Palatini curvature/torsion algebra),

$$\delta(*F)_{\mu\nu} = (*\delta F)_{\mu\nu} + \frac{1}{2} h(*F)_{\mu\nu} - \frac{1}{2} (h_\mu^\rho (*F)_{\rho\nu} + h_\nu^\rho (*F)_{\mu\rho}). \quad (\text{A10})$$

$[\mathcal{PT}, *] = 0$.

Because the metric is PT -even and the chosen orientation is preserved (A2), the Hodge map built from $(g, \epsilon_{\mu\nu\rho\sigma})$ commutes with PT :

$$\mathcal{PT}[*\alpha] = *\mathcal{PT}[\alpha] \quad \text{for any form } \alpha. \quad (\text{A11})$$

This identity is used both in the selection rules and in the projector proofs that involve p -form duals.

Appendix A.5. Boundary/Topology Posture and Improvement Currents

Assumption A4 is realized in either of the following equivalent ways:

- (i) *Compact, PT-invariant domains* with vanishing boundary flux: for any improvement current J^μ arising from integration by parts, $\int_{\partial\mathcal{D}} d\Sigma_\mu J^\mu = 0$.
- (ii) *Standard fall-offs* on asymptotically flat or spatially flat FRW patches, for which $\int d^4x \nabla_\mu J^\mu$ reduces to a surface integral that vanishes in the $R \rightarrow \infty$ limit. A sufficient set is

$$h_{\mu\nu} = \mathcal{O}(r^{-1-\delta}), \quad K^\lambda{}_{\mu\nu} = \mathcal{O}(r^{-2-\delta}), \quad \partial_\mu \epsilon = \mathcal{O}(r^{-\delta}), \quad \delta > 0, \quad (\text{A12})$$

which ensures $J^r = \mathcal{O}(r^{-2-\delta})$ so that the flux through a sphere of radius R decays as $R^{-\delta}$.

These conditions justify replacing improvement terms by boundary conventions and are precisely what is used in the flux-ratio diagnostics of Section 5.2.

Appendix A.6. Consequences Used in the Main Text

(C1) Palatini block-diagonalization.

Theorem A1 (A3) allows us to *project then vary* in the Palatini equations, so that the connection variation is algebraic and block-diagonal in the torsion irreps. Together with the selection rules (Prop. A2) this yields $S^\mu = q_{\lambda\mu\nu} = 0$ and the pure-trace map quoted in Section 4.

(C2) Route equivalence modulo boundary.

Self-adjointness (Prop. A1) and the boundary posture (Section A.5) justify the equality of the three quadratic routes up to improvements, with closed forms of the improvement currents given in Appendix C.

(C3) Equal-coefficient identity and $c_T=1$.

The star-variation identities (A9)–(A10) are used inside the ADM expansion behind the equal-coefficient identity $K - G = \partial_\mu (a^{-3} J_\Delta^\mu a^3)$ proven in Appendix D. The boundary posture then enforces $K=G$ and $c_T=1$ at quadratic order.

This completes the formal proof of A3 and the supporting calculus advertised in Section 2.

Appendix B. Irrep Projectors & No-Go for $q_{\lambda\mu\nu}(v)$

This appendix collects the group-theoretic ingredients used in Section 4: (i) the irreducible decomposition of the torsion tensor under the Lorentz group, (ii) explicit, idempotent projectors onto the trace, axial, and traceless sectors, (iii) the quadratic identity for $T^A{}_{BC} T_A{}^{BC}$ in our conventions, and (iv) the *single-vector no-go* that underlies the statement quoted in the main text as “Proposition B.1” for the $q_{\lambda\mu\nu}$ irrep. All statements are purely algebraic and hold before/after applying the scalar projector $\Pi_{\text{PT}}[\cdot]$; after projection all scalar contractions are real (Section 2).

Appendix B.1. Torsion as a Lorentz Representation and Its Algebra

In index language (spacetime indices), torsion is a rank-3 tensor antisymmetric in its last two indices, $T_{\lambda\mu\nu} = -T_{\lambda\nu\mu}$, with $4 \times 6 = 24$ independent components in $d=4$. The Lorentz-covariant irreducible content splits into

$$24 \simeq \underbrace{4}_{\text{trace } T_\mu} \oplus \underbrace{4}_{\text{axial } S^\mu} \oplus \underbrace{16}_{\text{traceless tensor } q_{\lambda\mu\nu}}, \quad \begin{cases} T_\mu \equiv T^\nu{}_{\mu\nu}, \\ S^\mu \equiv \epsilon^{\mu\nu\rho\sigma} T_{\nu\rho\sigma}, \\ q^\nu{}_{\mu\nu} = q_{\lambda\mu}{}^\mu = \epsilon^{\mu\nu\rho\sigma} q_{\nu\rho\sigma} = 0. \end{cases} \quad (\text{A13})$$

We use $\epsilon^{0123} = +1$ and the metric signature $(-, +, +, +)$, and we adopt the standard scalar product $(X, Y) \equiv X_{\lambda\mu\nu} Y^{\lambda\mu\nu}$ on this space.⁶

Appendix B.2. Idempotent Projectors

Define three linear maps $P^{(T)}, P^{(S)}, P^{(q)}$ on the torsion space by

$$(P^{(T)}T)_{\lambda\mu\nu} \equiv \frac{1}{3}(g_{\lambda\mu}T_\nu - g_{\lambda\nu}T_\mu), \quad (\text{A14})$$

$$(P^{(S)}T)_{\lambda\mu\nu} \equiv -\frac{1}{6}\epsilon_{\lambda\mu\nu\rho}S^\rho = -\frac{1}{6}\epsilon_{\lambda\mu\nu\rho}\epsilon^{\rho\alpha\beta\gamma}T_{\alpha\beta\gamma}, \quad (\text{A15})$$

$$(P^{(q)}T)_{\lambda\mu\nu} \equiv T_{\lambda\mu\nu} - (P^{(T)}T)_{\lambda\mu\nu} - (P^{(S)}T)_{\lambda\mu\nu}. \quad (\text{A16})$$

These are the unique Lorentz-covariant, algebraic (derivative-free) projectors onto the three irreps in Equation (A13). A direct computation shows:

$$(P^{(X)})^2 = P^{(X)}, \quad P^{(X)}P^{(Y)} = 0 \ (X \neq Y), \quad P^{(T)} + P^{(S)} + P^{(q)} = \mathbf{1}, \quad (\text{A17})$$

and the images obey by construction

$$\begin{aligned} (P^{(T)}T)^\nu_{\mu\nu} &= T_\mu, & \epsilon^{\mu\nu\rho\sigma}(P^{(T)}T)_{\nu\rho\sigma} &= 0, \\ (P^{(S)}T)^\nu_{\mu\nu} &= 0, & \epsilon^{\mu\nu\rho\sigma}(P^{(S)}T)_{\nu\rho\sigma} &= S^\mu, \\ (P^{(q)}T)^\nu_{\mu\nu} &= 0, & \epsilon^{\mu\nu\rho\sigma}(P^{(q)}T)_{\nu\rho\sigma} &= 0. \end{aligned} \quad (\text{A18})$$

Orthogonality and quadratic split.

With the scalar product $(\cdot, \cdot)$,

$$(P^{(X)}T, P^{(Y)}T) = 0 \ (X \neq Y), \quad (T, T) = (P^{(T)}T, P^{(T)}T) + (P^{(S)}T, P^{(S)}T) + (P^{(q)}T, P^{(q)}T). \quad (\text{A19})$$

Using (A14)–(A16) one finds the standard quadratic identity

$$T^A_{BC}T_A{}^{BC} = \frac{2}{3}T_\mu T^\mu + \frac{1}{24}S_\mu S^\mu + \tilde{q}_{\lambda\mu\nu}\tilde{q}^{\lambda\mu\nu}, \quad (\text{A20})$$

where we have denoted $\tilde{q} \equiv P^{(q)}T$ for the *standard* normalization of the traceless piece.

Normalization used in the main text.

For later convenience—and to match the coefficient choice used in Section 4—we rescale the traceless irrep by a constant factor and *define*

$$q_{\lambda\mu\nu} \equiv \frac{1}{\sqrt{2}}\tilde{q}_{\lambda\mu\nu}, \quad \implies \quad T^A_{BC}T_A{}^{BC} = \frac{2}{3}T_\mu T^\mu + \frac{1}{24}S_\mu S^\mu + 2q_{\lambda\mu\nu}q^{\lambda\mu\nu}. \quad (\text{A21})$$

The projector formulas (A14)–(A16) are unchanged; only the bookkeeping name “ q ” for the traceless image carries the fixed $\sqrt{2}$ factor.⁷ After applying $\Pi_{\text{PT}}[\cdot]$ to either side of (A21), the scalar is manifestly real (Thm. 1).

⁶ Frame and spacetime contractions are equivalent once $e^A{}_\mu$ is inserted; all formulas below hold with either $T^A{}_{BC}$ or $T_{\lambda\mu\nu}$, and we freely move between the two notations.

⁷ This harmless convention matches Equation (4.13) of the main text (positivity split) and simplifies a few numerical diagnostics; no physics depends on it.

Appendix B.3. Compatibility with the Scalar Projector

The projectors $P^{(X)}$ are algebraic and commute with $\Pi_{PT}[\cdot]$ at the scalar level: for any two torsions T, T' ,

$$\Pi_{PT}[(P^{(X)}T)_{\lambda\mu\nu} (P^{(Y)}T')^{\lambda\mu\nu}] = \delta^{XY} \Pi_{PT}[(P^{(X)}T)_{\lambda\mu\nu} (P^{(X)}T')^{\lambda\mu\nu}]. \quad (A22)$$

Moreover, the mixed scalar $T_\mu S^\mu$ is PT -odd and is annihilated by $\Pi_{PT}[\cdot]$ (Section 1). Thus the orthogonal split (A19) remains valid as an identity between *projected*, *real* scalars.

Appendix B.4. Proposition B.1: Single-Vector No-Go for the Traceless Irrep

We now formalize the statement invoked in Section 4.1.

Proposition A3 (single-vector no-go). *Let v_μ be any nonzero covector. There is no nonvanishing tensor of the form $q_{\lambda\mu\nu}(v)$, linear in v_μ , that (i) is antisymmetric in $\mu \leftrightarrow \nu$, (ii) obeys $q^\nu_{\mu\nu} = q_{\lambda\mu}{}^\mu = 0$, and (iii) satisfies $\epsilon^{\mu\nu\rho\sigma} q_{\nu\rho\sigma} = 0$. Equivalently, the traceless irrep cannot be constructed from a single vector.*

Proof. The most general Lorentz-covariant tensor built linearly from a single v_μ and antisymmetric in its last two indices is a linear combination of the two rank-one seeds

$$Y^A{}_{BC}(v) = \alpha \delta^A{}_{[B} v_{C]} + \beta \epsilon^A{}_{BCD} v^D, \quad (A23)$$

with real α, β . Compute its traces and axial contraction:

$$Y^\nu{}_{\mu\nu}(v) = \alpha \delta^\nu{}_{[\mu} v_{\nu]} = \frac{\alpha}{2} (\delta^\nu{}_\mu v_\nu - \delta^\nu{}_\nu v_\mu) = -\frac{3\alpha}{2} v_\mu, \quad (A24)$$

$$Y_{\lambda\mu}{}^\mu(v) = \alpha \delta_{\lambda[\mu} v_{\nu]} g^{\mu\nu} = 0, \quad (A25)$$

$$\epsilon^{\mu\nu\rho\sigma} Y_{\nu\rho\sigma}(v) = \beta \epsilon^{\mu\nu\rho\sigma} \epsilon_{\nu\rho\sigma D} v^D = -3! \beta \delta^\mu{}_D v^D = -6\beta v^\mu, \quad (A26)$$

where we used $\epsilon^{\mu\nu\rho\sigma} \epsilon_{\alpha\nu\rho\sigma} = -3! \delta^\mu{}_\alpha$ and $d=4$. Requiring the trace constraints $Y^\nu{}_{\mu\nu} = Y_{\lambda\mu}{}^\mu = 0$ forces $\alpha = 0$, and the axial constraint forces $\beta = 0$. Therefore the only admissible linear combination is the trivial one, $Y^A{}_{BC} \equiv 0$, proving the claim. $\square$

Corollary A1. *For any single covector v_μ the $P^{(q)}$ projector annihilates the two rank-one seeds: $P^{(q)}[\delta^A{}_{[B} v_{C]}] = 0 = P^{(q)}[\epsilon^A{}_{BCD} v^D]$.*

Appendix B.5. Consequences for the C1 Ansatz

Applying Prop. A3 to the most general *linear* ansatz with one derivative (Section 4.1),

$$T^A{}_{BC} = A \delta^A{}_{[B} \partial_{C]} \epsilon + B \epsilon^A{}_{BCD} \partial^D \epsilon + (\mathcal{P} \cdot \partial \epsilon)^A{}_{BC}, \quad (A27)$$

shows that the attempted traceless piece $(\mathcal{P} \cdot \partial \epsilon)^A{}_{BC}$ necessarily vanishes: it is a linear, single-vector construct and is thus killed by $P^{(q)}$ (Cor. A1). The ansatz collapses to

$$T^A{}_{BC} = A \delta^A{}_{[B} \partial_{C]} \epsilon + B \epsilon^A{}_{BCD} \partial^D \epsilon, \quad (A28)$$

recovering Equation (4.2) of the main text. The scalar projector $\Pi_{PT}[\cdot]$ removes PT -odd *scalars* built from the axial seed (Section 1), and the Palatini connection equation then sets $B=0$ while fixing $A = 2\eta$ under the trace lock $T_\mu = 3\eta \partial_\mu \epsilon$ (Section 4.1). This yields the uniqueness map $T^A{}_{BC} = 2\eta \delta^A{}_{[B} \partial_{C]} \epsilon$ quoted in Theorem 2.

Appendix B.6. Consistency Check with the Quadratic Invariant

With the normalization (A21) and the C1 map (so $S^\mu=0$ and $q_{\lambda\mu\nu}=0$),

$$\begin{aligned}\Pi_{\text{PT}}[T^A{}_{BC}T_A{}^{BC}] &= \frac{2}{3} T_\mu T^\mu = \frac{2}{3} (3\eta)^2 \Sigma_\epsilon = 6\eta^2 \Sigma_\epsilon, \\ I_T &\equiv -\frac{1}{4} \Pi_{\text{PT}}[T^A{}_{BC}T_A{}^{BC}] = -6\eta^2 \Sigma_\epsilon.\end{aligned}\quad (\text{A29})$$

as used throughout Secs. 4–5. Here $\Sigma_\epsilon \equiv \Pi_{\text{PT}}[(\partial\epsilon)^2]$ is the *projected, real* scalar, and the sign bookkeeping is carried by $\text{sgn}(\Sigma_\epsilon) \equiv \text{sgn}(\Sigma_\epsilon)$; the unit 1-form $\hat{n}_\mu \equiv \partial_\mu \epsilon / \sqrt{|\Sigma_\epsilon|}$, the canonical traceless rank-one matrix $\mathcal{T}^\mu{}_\nu \equiv \text{sgn}(\Sigma_\epsilon)(\hat{n}^\mu \hat{n}_\nu - \frac{1}{4} \delta^\mu{}_\nu)$, and the trace scale $\tau \equiv \hat{n}^\mu T_\mu = 3\eta \text{sgn}(\Sigma_\epsilon) \sqrt{|\Sigma_\epsilon|}$ are recalled from Section 3.

Appendix B.7. Edge Cases and Patches

On loci where $\Sigma_\epsilon = 0$ the normalized direction $\hat{n}$ is defined patchwise (or by continuity); all algebraic projector statements remain valid, and the conclusions above hold on any patch with $\Sigma_\epsilon \neq 0$. Global/topological subtleties (multi-valued ϵ , nontrivial bundles) lie outside the posture A1–A5 (Section 2).

Summary of Appendix B. We have given explicit, idempotent projectors onto the three torsion irreps, fixed the quadratic identity in the normalization used in the main text, and proved the *single-vector no-go*: from one covector v_μ no nonzero traceless torsion irrep can be built. This reduces the most general one-derivative ansatz to the $\{\delta^A{}_{[B}\partial_{C]}\epsilon, \epsilon^A{}_{BCD}\partial^D\epsilon\}$ span, after which the Palatini equations and the scalar projector select the pure-trace map used in the C1 uniqueness theorem.

Appendix C. Appendix C: Three–Chain Reductions & Improvement Currents (σ_ϵ Scheme)

Scope (σ_ϵ scheme and naming). This appendix provides the paper-checkable reductions behind Section 5 under the *sign-compensated* convention $\text{sgn}(\Sigma_\epsilon) I_T \equiv \text{sgn}(\Sigma_\epsilon) I_T = -6\eta^2 |\Sigma_\epsilon|$ and the rank-one determinant route naming (formerly “DBI”-type; not Born–Infeld gravity): (i) a rank-one determinant route built out of the canonical traceless matrix $\mathcal{T}^\mu{}_\nu$, (ii) a closed–metric rank-one deformation, and (iii) the *PT*–even CS/Nieh–Yan shadow. At quadratic order, *each route* reduces in the bulk to the same invariant line

$$\begin{aligned}\delta^2 \mathcal{L}_X &= A_\star \sqrt{-g} \text{sgn}(\Sigma_\epsilon) I_T + \nabla_\mu J_X^\mu, \\ A_\star &= \lambda^2/8, \quad X \in \{\text{rank-one determinant route}, \text{CM}, \text{CS}^+\}.\end{aligned}\quad (\text{A30})$$

with improvements $\nabla_\mu J_X^\mu$ differing by boundary choices (A4/A5). Closed representatives for J^μ are given on FRW/weak-field backgrounds.

Notation and key relation. We use the preamble shorthands $\Sigma_\epsilon \equiv \Pi_{\text{PT}}[(\partial\epsilon)^2]$, $\text{sgn}(\Sigma_\epsilon) \equiv \text{sgn}(\Sigma_\epsilon)$, $\hat{n}_\mu \equiv \frac{\partial_\mu \epsilon}{\sqrt{|\Sigma_\epsilon|}}$, $\mathcal{T}^\mu{}_\nu \equiv \text{sgn}(\Sigma_\epsilon)(\hat{n}^\mu \hat{n}_\nu - \frac{1}{4} \delta^\mu{}_\nu)$, $\tau \equiv \hat{n}^\mu T_\mu = 3\eta \text{sgn}(\Sigma_\epsilon) \sqrt{|\Sigma_\epsilon|}$ so that $\text{Tr} \mathcal{T} = 0$, $\text{Tr}(\mathcal{T}^2) = \frac{3}{4}$, and (after C1)

$$\tau^2 = 9\eta^2 |\Sigma_\epsilon| = \frac{3}{2} (-\text{sgn}(\Sigma_\epsilon) I_T), \quad \text{sgn}(\Sigma_\epsilon) I_T \equiv \text{sgn}(\Sigma_\epsilon) I_T = -6\eta^2 |\Sigma_\epsilon|.$$

Appendix C.1. rank-one determinant route: Determinant Algebra with the $\frac{2}{3}$ Normalization

Consider the rank-one determinant route Lagrangian

$$\mathcal{L}_{\text{ROD}} = \sqrt{-g} \left(\sqrt{\det[\mathbf{1} + \frac{2}{3} \lambda \tau \mathcal{T}]} - 1 \right), \quad \lambda \in \mathbb{R}. \quad (\text{A31})$$

Using $\sqrt{\det(\mathbf{1} + X)} = 1 + \frac{1}{2}\text{Tr}X + \frac{1}{8}[(\text{Tr}X)^2 - 2\text{Tr}X^2] + \mathcal{O}(X^3)$ and $\text{Tr}\mathcal{T} = 0$, the quadratic piece is

$$\begin{aligned}\delta^2 \mathcal{L}_{\text{ROD}} &= \sqrt{-g} \frac{1}{8} \left[-2 \text{Tr} \left(\left(\frac{2}{3} \lambda \tau \right)^2 \mathcal{T}^2 \right) \right] = -\frac{1}{4} \left(\frac{2}{3} \right)^2 \lambda^2 \sqrt{-g} \tau^2 \text{Tr}(\mathcal{T}^2) \\ &= -\frac{1}{4} \cdot \frac{4}{9} \cdot \frac{3}{4} \lambda^2 \sqrt{-g} \tau^2 = -\frac{1}{12} \lambda^2 \sqrt{-g} \tau^2.\end{aligned}\quad (\text{A32})$$

With $\tau^2 = 9\eta^2 |\Sigma_\epsilon|$ and $\text{sgn}(\Sigma_\epsilon) I_T = -6\eta^2 |\Sigma_\epsilon|$,

$$-\frac{1}{12} \lambda^2 \tau^2 = -\frac{1}{12} \lambda^2 (9\eta^2 |\Sigma_\epsilon|) = -\frac{3}{4} \lambda^2 \eta^2 |\Sigma_\epsilon| = \left(\frac{\lambda^2}{8} \right) \text{sgn}(\Sigma_\epsilon) I_T,$$

so that

$$\boxed{\delta^2 \mathcal{L}_{\text{ROD}} = A_\star \sqrt{-g} \text{sgn}(\Sigma_\epsilon) I_T + \nabla_\mu J_{\text{ROD}}^\mu, \quad A_\star = \frac{\lambda^2}{8}.} \quad (\text{A33})$$

At the bulk-density level one may take $J_{\text{ROD}}^\mu \equiv 0$. For unified boundary diagnostics we adopt a common canonical representative J_{can}^μ for *all* three routes (Section C.4).

Appendix C.2. Closed–Metric Route: Rank–One Deformation Equals rank-one determinant route to $\mathcal{O}(\mathcal{T}^2)$

Take the rank–one deformation

$$\tilde{g}_{\mu\nu} = g_{\mu\nu} + \frac{2}{3} \lambda \tau \hat{n}_\mu \hat{n}_\nu,$$

so that

$$\sqrt{-\tilde{g}} = \sqrt{-g} \sqrt{\det \left[\mathbf{1} + \frac{2}{3} \lambda \tau \mathcal{T} \right]}.$$

Define

$$\mathcal{L}_{\text{CM}} \equiv \sqrt{-\tilde{g}} - \sqrt{-g}.$$

The same algebra gives

$$\boxed{\delta^2 \mathcal{L}_{\text{CM}} = A_\star \sqrt{-g} \text{sgn}(\Sigma_\epsilon) I_T + \nabla_\mu J_{\text{CM}}^\mu, \quad A_\star = \frac{\lambda^2}{8}.} \quad (\text{A34})$$

Thus rank-one determinant route and CM have the *same bulk* coefficient $A_\star$ and differ only by improvements.

Appendix C.3. PT–Even CS/Nieh–Yan Shadow: Quadratic Reduction (σ_ϵ)

Using $T^A \wedge T_A = d(e^A \wedge T_A) - e^A \wedge e^B \wedge R_{AB}$, applying $*$ and the scalar PT projector (A2), and evaluating after C1, the PT –even piece reduces at quadratic order to

$$\boxed{\delta^2 \mathcal{L}_{\text{CS}^+} = A_\star \sqrt{-g} \text{sgn}(\Sigma_\epsilon) I_T + \nabla_\mu J_{\text{CS}^+}^\mu, \quad A_\star = \frac{\lambda^2}{8},} \quad (\text{A35})$$

where the Nieh–Yan assignment (A5) reshuffles only boundary conventions inside the PT –even sector.

Appendix C.4. A Universal Canonical Improvement at Quadratic Order

For route–by–route flux comparisons it is convenient to select the *same* improvement representative for all routes:

$$\boxed{J_{\text{can}}^\mu[h^{\text{TT}}] \equiv \frac{A_\star}{2} a^3 \left(h_{ij}^{\text{TT}} \nabla^\mu h_{ij}^{\text{TT}} - \delta^\mu_0 h_{ij}^{\text{TT}} \dot{h}_{ij}^{\text{TT}} \right),} \quad (\text{A36})$$

so that we adopt the convention

$$\delta^2 \mathcal{L}_X = A_\star \sqrt{-g} \text{sgn}(\Sigma_\epsilon) I_T + \nabla_\mu J_{\text{can}}^\mu, \quad X \in \{\text{rank} - \text{onedeterminantroute}, \text{CM}, \text{CS}^+\}. \quad (\text{A37})$$

Different representatives differ by $\nabla_\nu X^{[\mu\nu]}$ and yield identical integrated fluxes under A4.

Check (FRW).

On spatially flat FRW in TT gauge,

$$\nabla_\mu J_{\text{can}}^\mu = \frac{A_\star a^3}{2} \left(h_{ij}^{\text{TT}} \dot{h}_{ij}^{\text{TT}} - \frac{\partial_k h_{ij}^{\text{TT}} \partial_k h_{ij}^{\text{TT}}}{a^2} \right) + \partial_\tau \left(\frac{3}{2} A_\star a^2 \dot{a} h_{ij}^{\text{TT}} h_{ij}^{\text{TT}} \right), \quad (\text{A38})$$

i.e. the canonical reshuffling between TT kinetic/gradient bilinears plus a pure time boundary term that integrates to zero with A4 fall-offs.

Appendix C.5. Closed Forms for FRW and Weak Field

FRW.

For $ds^2 = a(\tau)^2(-d\tau^2 + dx^2)$ and homogeneous $\epsilon(\tau)$,

$$J_{\text{can}}^0 = \frac{A_\star}{2} a^3 h_{ij}^{\text{TT}} \dot{h}_{ij}^{\text{TT}}, \quad J_{\text{can}}^i = \frac{A_\star}{2} a h_{jk}^{\text{TT}} \partial^i h_{jk}^{\text{TT}}, \quad (\text{A39})$$

which satisfies (A38).

Weak field (AF).

At $a \rightarrow 1$,

$$J_{\text{can}}^0 = \frac{A_\star}{2} h_{ij}^{\text{TT}} \dot{h}_{ij}^{\text{TT}}, \quad J_{\text{can}}^i = \frac{A_\star}{2} h_{jk}^{\text{TT}} \partial^i h_{jk}^{\text{TT}}. \quad (\text{A40})$$

Appendix C.6. Flux–Ratio Identity & Finite–Domain Convergence

With the unified choice (A37), $\int_{\partial\mathcal{D}} d\Sigma_\mu J_{\text{ROD}}^\mu = \int_{\partial\mathcal{D}} d\Sigma_\mu J_{\text{CM}}^\mu = \int_{\partial\mathcal{D}} d\Sigma_\mu J_{\text{CS}^+}^\mu$, hence

$$\mathcal{R}_{X/Y}[\partial\mathcal{D}] \equiv \frac{\int_{\partial\mathcal{D}} d\Sigma_\mu J_X^\mu}{\int_{\partial\mathcal{D}} d\Sigma_\mu J_Y^\mu} = 1, \quad X, Y \in \{\text{rank-one determinant route, CM, CS}^+\}. \quad (\text{A41})$$

On finite FRW balls (or AF shells) residuals scale away with the radius R , agreeing with Sec. V.

Summary of Appendix C

At quadratic order and under A1–A5 plus C1, the rank-one determinant route, closed–metric, and PT –even CS/Nieh–Yan routes share the same bulk reduction $\delta^2 \mathcal{L}_X = A_\star \sqrt{-g} \text{sgn}(\Sigma_\epsilon) I_T \pmod{\nabla_\mu J^\mu}$, $A_\star = \lambda^2/8$. A single canonical improvement J_{can}^μ (Eqs. (A36)–(A38)) is used for all routes and underlies the flux–ratio plots in Sec. V. [nb : fig_c2_coeff_compare.py; fig_flux_ratio.py]

Appendix D. Appendix D: Mixing Matrix and the Equal–Coefficient Identity

This appendix contains (i) extraction rules and tables for the 2×2 mixing matrix used in Section 6, including a *non-collinearity* proof of its two row vectors on admissible backgrounds, and (ii) a covariant derivation of the *equal-coefficient identity* $K(w) - G(w) = a^{-3} \partial_\mu (a^3 J_\Delta^\mu)$ quoted in Section 7. We assume A1–A5, the scalar PT projector (Section 2), and the C1 map $T_\mu = 3\eta \partial_\mu \epsilon$. Projected scalars are real by construction; the σ_ϵ scheme only affects the bulk line through $\text{sgn}(\Sigma_\epsilon) I_T = \text{sgn}(\Sigma_\epsilon) I_T = -6\eta^2 |\Sigma_\epsilon|$, not the kinematical identity $K - G = \partial J$.

Variational domain (used below).

We take variations with compact support on spatial slices or with FRW/AF fall-offs: $h_{ij}^{\text{TT}} = O(r^{-1-\sigma})$, $\delta N = O(r^{-2-\sigma})$, $\partial_i N^i = O(r^{-2-\sigma})$ with $\sigma > 0$. Then $J_\Delta^\mu = O(r^{-3-\sigma})$ and $\delta \int d^4x a^3 \partial_\mu J_\Delta^\mu = \int d\Sigma_\mu a^3 \delta J_\Delta^\mu = 0$.

Appendix D.1. ADM Conventions and Extraction of Mixing Entries

With $ds^2 = -N^2 d\tau^2 + \gamma_{ij}(dx^i + N^i d\tau)(dx^j + N^j d\tau)$, $N = 1 + \delta N$, $N_i = \partial_i \chi + N_i^T$ ($\partial_i N_i^T = 0$), and $\gamma_{ij} = a^2(\tau)(\delta_{ij} + h_{ij})$, we project $h_{ij} \rightarrow h_{ij}^{TT}$ with $\partial_i h_{ij}^{TT} = 0 = \delta^{ij} h_{ij}^{TT}$. The quadratic Lagrangian takes the block form

$$\delta^2 \mathcal{L} = \frac{M_{\text{Pl}}^2}{8} a^3 \left[K(w) \dot{h}_{ij}^{TT} \dot{h}_{ij}^{TT} - G(w) \frac{(\partial_k h_{ij}^{TT})^2}{a^2} \right] + a^3 (h^{TT} \cdot \mathbf{M}(w) \cdot \Phi) + \frac{a^3}{2} \Phi \cdot \mathbf{P} \cdot \Phi, \quad (\text{A42})$$

where $\Phi = \{\delta N, \partial_i N^i, \dots\}$ collects nonpropagating pieces. Define the *dimensionless* mixing entries by

$$[\mathbf{M}(w)]_{h^{TT}-\delta N} \equiv \frac{8}{M_{\text{Pl}}^2} \frac{1}{a^3} \frac{\partial}{\partial(\delta N)} \left(\delta^2 \mathcal{L} \right) \Big|_{\text{bilinear in } h^{TT}, \delta N}, \quad (\text{A43})$$

$$[\mathbf{M}(w)]_{h^{TT}-\partial_i N^i} \equiv \frac{8}{M_{\text{Pl}}^2} \frac{1}{a^3} \frac{\partial}{\partial(\partial_i N^i)} \left(\delta^2 \mathcal{L} \right) \Big|_{\text{bilinear in } h^{TT}, \partial_i N^i}. \quad (\text{A44})$$

For $\mathcal{L}_{\text{chain}}(w) = w_{\text{ROD}} \mathcal{L}_{\text{ROD}} + w_{\text{CM}} \mathcal{L}_{\text{CM}}$, the mixing block is linear in w and proportional to $\eta^2 \Sigma_\epsilon$:

$$\begin{aligned} [\mathbf{M}(w)]_{h^{TT}-\delta N} &= \left(\mu_{\text{ROD}}^{(N)} w_{\text{ROD}} + \mu_{\text{CM}}^{(N)} w_{\text{CM}} \right) \eta^2 \Sigma_\epsilon, \\ [\mathbf{M}(w)]_{h^{TT}-\partial_i N^i} &= \left(\mu_{\text{ROD}}^{(\nabla N)} w_{\text{ROD}} + \mu_{\text{CM}}^{(\nabla N)} w_{\text{CM}} \right) \eta^2 \Sigma_\epsilon, \end{aligned} \quad (\text{A45})$$

defining the four dimensionless coefficients $\mu_{\text{ROD}}^{(N)}, \mu_{\text{CM}}^{(N)}, \mu_{\text{ROD}}^{(\nabla N)}, \mu_{\text{CM}}^{(\nabla N)}$.

Appendix D.2. Background Invariants and Compact Parametrization

Introduce the two (dimensionful) background invariants

$$\mathcal{H} \equiv \frac{a'}{a}, \quad \mathcal{Y} \equiv \partial_\tau \ln(\tau/a) = \frac{\tau'}{\tau} - \mathcal{H}, \quad (\text{A46})$$

where prime denotes conformal-time derivative. Each μ admits a linear decomposition

$$\mu_X^{(Y)} = c_{X,0}^{(Y)} + c_{X,\mathcal{H}}^{(Y)} \mathcal{H} + c_{X,\mathcal{Y}}^{(Y)} \mathcal{Y}, \quad X \in \{\text{ROD}, \text{CM}\}, \quad Y \in \{N, \nabla N\}, \quad (\text{A47})$$

with c 's real $\mathcal{O}(1)$ numbers fixed by the quadratic expansion rules (route Jacobians plus contorsion under C1).

Appendix D.3. Coefficient Tables (FRW and Weak Field)

Table A1. FRW coefficients with $\epsilon = \epsilon(\tau)$. Entries are the dimensionless μ 's of (A45), written as $\mu = c_0 + c_{\mathcal{H}} \mathcal{H} + c_{\mathcal{Y}} \mathcal{Y}$ [Equation (A47)]. An overall factor $\eta^2 \Sigma_\epsilon$ multiplies the mixing (*not* shown here).

Coefficient	c_0	$c_{\mathcal{H}}$	$c_{\mathcal{Y}}$
$\mu_{\text{ROD}}^{(N)}$	$c_{\text{ROD},0}^{(N)}$	$c_{\text{ROD},\mathcal{H}}^{(N)}$	$c_{\text{ROD},\mathcal{Y}}^{(N)}$
$\mu_{\text{ROD}}^{(\nabla N)}$	$c_{\text{ROD},0}^{(\nabla N)}$	$c_{\text{ROD},\mathcal{H}}^{(\nabla N)}$	$c_{\text{ROD},\mathcal{Y}}^{(\nabla N)}$
$\mu_{\text{CM}}^{(N)}$	$c_{\text{CM},0}^{(N)}$	$c_{\text{CM},\mathcal{H}}^{(N)}$	$c_{\text{CM},\mathcal{Y}}^{(N)}$
$\mu_{\text{CM}}^{(\nabla N)}$	$c_{\text{CM},0}^{(\nabla N)}$	$c_{\text{CM},\mathcal{H}}^{(\nabla N)}$	$c_{\text{CM},\mathcal{Y}}^{(\nabla N)}$

Regeneration recipe (FRW). (1) Write $\delta g_{\mu\nu}$ in terms of $(\delta N, \partial_i N^i, h_{ij}^{TT})$. (2) Expand $\mathcal{T}, \tau$ to linear order using $\hat{n}_\mu = \partial_\mu \epsilon / \sqrt{|\Sigma_\epsilon|}$ and C1. (3) Insert into $\mathcal{L}_{\text{ROD}}$ and $\mathcal{L}_{\text{CM}}$ with the $\frac{2}{3}$ scaling. (4) Keep only $h^{TT} \delta N$ and $h^{TT} \partial_i N^i$ bilinears to read off $c_0, c_{\mathcal{H}}, c_{\mathcal{Y}}$.

Table A2. Weak field (AF) coefficients with slowly varying $\epsilon(t, \mathbf{x})$. Set $a \rightarrow 1$ ($\mathcal{H} = 0$) and define $\mathcal{Y} \equiv \partial_t \ln \tau$. Overall factor $\eta^2 \Sigma_\epsilon$ multiplies the mixing (not shown here).

Coefficient	c_0	$c_{\mathcal{Y}}$
$\mu_{\text{ROD}}^{(N)}$	$\tilde{c}_{\text{ROD},0}^{(N)}$	$\tilde{c}_{\text{ROD},\mathcal{Y}}^{(N)}$
$\mu_{\text{ROD}}^{(\nabla N)}$	$\tilde{c}_{\text{ROD},0}^{(\nabla N)}$	$\tilde{c}_{\text{ROD},\mathcal{Y}}^{(\nabla N)}$
$\mu_{\text{CM}}^{(N)}$	$\tilde{c}_{\text{CM},0}^{(N)}$	$\tilde{c}_{\text{CM},\mathcal{Y}}^{(N)}$
$\mu_{\text{CM}}^{(\nabla N)}$	$\tilde{c}_{\text{CM},0}^{(\nabla N)}$	$\tilde{c}_{\text{CM},\mathcal{Y}}^{(\nabla N)}$

Regeneration recipe (weak field). Repeat the FRW steps with $a \rightarrow 1$, replace conformal by physical time, and expand consistently in spatial gradients so that Σ_ϵ remains the projected, real scalar.

Closed-form labels (for Sec. VI cross-references).

For later citation we record the compact forms as Equations (D.12)–(D.15):

$$\mu_{\text{ROD}}^{(N)} = c_{\text{ROD},0}^{(N)} + c_{\text{ROD},\mathcal{H}}^{(N)} \mathcal{H} + c_{\text{ROD},\mathcal{Y}}^{(N)} \mathcal{Y}, \quad (\text{D.12})$$

$$\mu_{\text{CM}}^{(N)} = c_{\text{CM},0}^{(N)} + c_{\text{CM},\mathcal{H}}^{(N)} \mathcal{H} + c_{\text{CM},\mathcal{Y}}^{(N)} \mathcal{Y}, \quad (\text{D.13})$$

$$\mu_{\text{ROD}}^{(\nabla N)} = c_{\text{ROD},0}^{(\nabla N)} + c_{\text{ROD},\mathcal{H}}^{(\nabla N)} \mathcal{H} + c_{\text{ROD},\mathcal{Y}}^{(\nabla N)} \mathcal{Y}, \quad (\text{D.14})$$

$$\mu_{\text{CM}}^{(\nabla N)} = c_{\text{CM},0}^{(\nabla N)} + c_{\text{CM},\mathcal{H}}^{(\nabla N)} \mathcal{H} + c_{\text{CM},\mathcal{Y}}^{(\nabla N)} \mathcal{Y}. \quad (\text{D.15})$$

Appendix D.4. Non-Collinearity of the Two Locking Equations

Let the two rows be $\mathbf{r}_1 = (\mu_{\text{ROD}}^{(N)}, \mu_{\text{CM}}^{(N)})$, $\mathbf{r}_2 = (\mu_{\text{ROD}}^{(\nabla N)}, \mu_{\text{CM}}^{(\nabla N)})$.

Proposition A4 (Non-collinearity). *Under A1–A5, C1, and the one-derivative-per-building-block posture, $\mathbf{r}_1 \nparallel \mathbf{r}_2$ on any admissible background with at least one of $\{\mathcal{H}, \mathcal{Y}\} \neq 0$ and $\nabla_\mu \hat{n}_\nu \neq 0$. Equivalently,*

$$\Delta \equiv \mu_{\text{ROD}}^{(N)} \mu_{\text{CM}}^{(\nabla N)} - \mu_{\text{ROD}}^{(\nabla N)} \mu_{\text{CM}}^{(N)} \neq 0.$$

Proof (sketch). Using (A47), proportionality would require $\rho_X \equiv \mu_X^{(N)} / \mu_X^{(\nabla N)}$ to be route-independent and constant under independent shifts of $\mathcal{H}$ and $\mathcal{Y}$. But $\partial_{\mathcal{H}} \rho_X$ and $\partial_{\mathcal{Y}} \rho_X$ cannot both vanish for $X = \text{ROD, CM}$ unless $c_{X,\mathcal{H}}^{(Y)} = c_{X,\mathcal{Y}}^{(Y)} = 0$ (static trivial branch $\mathcal{H}=\mathcal{Y}=0$, $\nabla \hat{n} = 0$). Hence $\mathbf{r}_1 \nparallel \mathbf{r}_2$ generically. $\square$

Small- k structure (symbolic).

On weak-field patches one finds

$$\Delta(k, \Sigma_\epsilon) = \alpha_0 \eta^4 \Sigma_\epsilon^2 k^2 + O(k^4, \Sigma_\epsilon^3), \quad \alpha_0 \neq 0,$$

so rank loss occurs only on the measure-zero set $\{k = 0\} \cup \{\Sigma_\epsilon = 0\}$ or special foliations.

Appendix D.5. Equal-Coefficient Identity: Covariant Derivation and Gauge Shift

Define the quadratic current (indices contracted with the background spatial metric)

$$J_\Delta^\mu = \frac{M_{\text{Pl}}^2}{8} \begin{pmatrix} J_\Delta^0 = \Pi^{ij} h_{ij}^{\text{TT}} \\ J_\Delta^k = -\Xi^{ij} \partial_k h_{ij}^{\text{TT}} \end{pmatrix}, \quad (\text{A48})$$

with $\Pi^{ij} \equiv \delta(\delta^2 \mathcal{L}) / \delta h_{ij}^{\text{TT}}$ and $\Xi^{ij} \equiv \delta(\delta^2 \mathcal{L}) / \delta(\partial^2 h_{ij}^{\text{TT}})$. Using TT conditions, the rank-one traceless normalization, and the bulk equality of the routes (Sec. V),

$$K(w) - G(w) = \frac{1}{a^3} \partial_\mu (a^3 J_\Delta^\mu). \quad (\text{A49})$$

Under $\delta_{\xi} h_{ij}^{\text{TT}} = 0$, $\delta_{\xi} h_{0\mu} = \partial_{\mu} \xi_0 + \dots$, the representative shifts as $\delta_{\xi} J_{\Delta}^{\mu} = \nabla_{\nu} X^{[\mu\nu]}[\xi]$, hence the *integrated* identity is gauge independent.

FRW and weak-field representatives.

On spatially flat FRW ($ds^2 = a^2(-d\tau^2 + d\mathbf{x}^2)$, homogeneous ϵ),

$$J_{\Delta}^0 = \frac{A_{\star} \eta^2}{2} \Sigma_{\epsilon} h_{ij}^{\text{TT}} h_{ij}^{\text{TT}} + \nabla_i(\dots), \quad (\text{A50})$$

$$J_{\Delta}^k = -\frac{A_{\star} \eta^2}{2} \Sigma_{\epsilon} \frac{1}{a^2} h_{ij}^{\text{TT}} \partial_k h_{ij}^{\text{TT}} + \partial_{\ell}(\dots)^{[k\ell]}, \quad (\text{A51})$$

so $K - G = a^{-3} \partial_{\tau}(a^3 J_{\Delta}^0) + a^{-3} \partial_k(a^3 J_{\Delta}^k)$ is a total divergence. In weak field (AF, $a \rightarrow 1$),

$$J_{\Delta}^0 = \frac{A_{\star} \eta^2}{2} \Sigma_{\epsilon} h_{ij}^{\text{TT}} h_{ij}^{\text{TT}} + \partial_i(\dots)^i, \quad J_{\Delta}^k = -\frac{A_{\star} \eta^2}{2} \Sigma_{\epsilon} h_{ij}^{\text{TT}} \partial_k h_{ij}^{\text{TT}} + \partial_{\ell}(\dots)^{[k\ell]}.$$

Boundary consequence and luminosity.

With the variational domain above (A4), $\int a^3(K - G) = \int \partial_{\mu}(a^3 J_{\Delta}^{\mu}) = 0$. At the *locked* weights $w^{\star}$ (Sec. VI; no TT-nonTT mixing), $K(w^{\star}) = G(w^{\star}) \Rightarrow c_T^2 = 1$ at quadratic order, slicing-independently.

Appendix D.6. Locked Weights and Summary Box

Non-collinearity with the two mixing equations yields a unique ratio $w_{\text{ROD}}^{\star}/w_{\text{CM}}^{\star}$ (up to the GR normalization). At these weights, $K - G = a^{-3} \partial_{\mu}(a^3 J_{\Delta}^{\mu})$ implies exact luminosity.

Appendix D (at a glance).

- *Mixing matrix.* Four dimensionless coefficients $\mu_X^{(Y)}$ ($X = \text{ROD, CM}$; $Y = N, \nabla N$) control the TT-constraint mixing with an overall $\eta^2 \Sigma_{\epsilon}$ factor.
- *Non-collinearity.* The two locking rows are not proportional on evolving admissible backgrounds; the determinant Δ is generically nonzero ($\Delta \sim \alpha_0 \eta^4 \Sigma_{\epsilon}^2 k^2 + \dots$).
- *Equal-coefficient identity.* $K - G = a^{-3} \partial_{\mu}(a^3 J_{\Delta}^{\mu})$; under A4 this yields $K = G$ and $c_T^2 = 1$ at the locked weights $w^{\star}$.

Reproducibility Note

Tables A1–A2 come from a single symbolic pipeline that: (i) expands $\mathcal{T}$ and τ to linear order in ADM variables with the $\frac{2}{3}$ normalization, (ii) forms the rank-one determinant route/CM quadratic densities, (iii) projects onto the bilinears of Eqs. (A43)–(A44). Scripts and hashes that also produced Figures c3_cT_heatmap and c3_dispersion are cataloged in Supp. R. [nb :

mixing_matrix_extract.py]

Appendix E. Dirac Coupling, Field Redefinition, and the Nieh–Yan Scope

This appendix supplies the details promised in Section 8: (i) the torsion–fermion couplings in our conventions, (ii) the precise field redefinition that removes the trace channel once the pure-trace map (C1) is imposed, (iii) the role of the Nieh–Yan density as a boundary convention (A5), and (iv) the scope limitations on nontrivial topology (A4).

Throughout we work under the global posture A1–A5 of Section 2, use metric signature $\eta_{AB} = \text{diag}(-, +, +, +)$, and adopt the gamma-matrix conventions

$$\{\gamma^A, \gamma^B\} = 2\eta^{AB}, \quad \gamma_5 \equiv i\gamma^0\gamma^1\gamma^2\gamma^3, \quad \epsilon^{0123} = +1, \quad (\text{A52})$$

so that $\gamma^{AB} \equiv \frac{1}{2}[\gamma^A, \gamma^B]$ and $\gamma^{ABC} \equiv \gamma^{[A}\gamma^B\gamma^{C]} = i\epsilon^{ABCD}\gamma_5\gamma_D$.

Appendix E.1. Minimal Dirac Coupling in Riemann–Cartan Space

The minimally coupled Dirac Lagrangian is

$$\mathcal{L}_\psi = e \bar{\psi} \left(i e_A^\mu \gamma^A D_\mu - m \right) \psi, \quad D_\mu \psi = \partial_\mu \psi + \frac{1}{4} \omega_{\mu}^{AB} \gamma_{AB} \psi, \quad e \equiv \sqrt{-g}. \quad (\text{A53})$$

Splitting the spin connection into Levi–Civita plus contorsion, $\omega_{\mu}^{AB} = \omega_{\mu}^{AB}(\text{LC}) + K_{\mu}^{AB}$, the *torsion-dependent* piece of (A53) is

$$\Delta \mathcal{L}_\psi[K] = \frac{i}{4} e \bar{\psi} e_A^\mu \gamma^A K_{\mu}^{BC} \gamma_{BC} \psi = \frac{i}{4} e \bar{\psi} \gamma^A K_{BCA} \gamma^{BC} \psi, \quad (\text{A54})$$

where tangent indices are moved with η_{AB} , and $K_{ABC} \equiv e_A^\mu K_{BC\mu}$.

It is convenient to decompose contorsion into the standard torsion irreps (vector trace T_A , axial S^A , and traceless q_{ABC}):

$$K_{ABC} = \frac{1}{3} (\eta_{AC} T_B - \eta_{AB} T_C) - \frac{1}{6} \epsilon_{ABCD} S^D + q_{ABC}, \quad T_A \equiv T^B_{AB}, \quad S^A \equiv \epsilon^{ABCD} T_{BCD}, \quad (\text{A55})$$

with $q^B_{AB} = 0$ and $\epsilon^{ABCD} q_{BCD} = 0$. Using the gamma identities listed below Equation (A52) and the antisymmetry of K_{BCA} in (B, C) , Equation (A54) reduces to

$$\Delta \mathcal{L}_\psi[K] = e \left(c_A S_\mu J_5^\mu + c_V T_\mu J^\mu \right) + e C^{\mu\nu\rho}(q) \bar{\psi} \gamma_{[\mu} \gamma_\nu \gamma_{\rho]} \psi, \quad (\text{A56})$$

where $J^\mu \equiv \bar{\psi} \gamma^\mu \psi$ and $J_5^\mu \equiv \bar{\psi} \gamma^\mu \gamma_5 \psi$. In our normalization (A52) the *numerical* coefficients are

$$\boxed{c_A = \frac{3}{8}, \quad c_V = -\frac{1}{4}} \quad (\text{A57})$$

while the q -channel is proportional to the totally antisymmetric part of q , which vanishes by definition; equivalently, the q -coupling can be re-expressed in terms of S^μ and drops out when $S^\mu = 0$. Therefore the torsion-induced Dirac channels in our conventions are precisely

$$\boxed{\Delta \mathcal{L}_\psi[T, S] = e \left(\frac{3}{8} S_\mu J_5^\mu - \frac{1}{4} T_\mu J^\mu \right)}, \quad (\text{A58})$$

up to PT -even improvement terms annihilated by A1/A4 in the bulk.

Cross-check.

Equation (A58) reproduces the two-channel bookkeeping used in Section 8 with $c_A = \frac{3}{8}$ and $c_V = -\frac{1}{4}$. Any alternative gamma/Hodge convention simply rescales (A57) by an overall sign; our later conclusions (vanishing axial channel by C1 and removability of the trace channel) are insensitive to such simultaneous flips.

Appendix E.2. C1 Implies no Axial Channel; the Trace Channel is Removable

By the uniqueness theorem (C1), Equation (11), the axial and traceless torsion irreps vanish, $S^\mu = 0$, $q_{\lambda\mu\nu} = 0$, and the trace aligns with the spurion gradient, $T_\mu = 3\eta \partial_\mu \epsilon$. Equation (A58) therefore reduces to the single trace channel

$$\mathcal{L}_{\psi, T}^{(\text{trace})} = e c_V T_\mu J^\mu = 3c_V \eta e (\partial_\mu \epsilon) J^\mu. \quad (\text{A59})$$

Proposition A5 (Vector rephasing removes the trace channel). *Consider the local vector phase redefinition $\psi \rightarrow \psi' = e^{i\alpha\epsilon} \psi$, $\bar{\psi} \rightarrow \bar{\psi}' = \bar{\psi} e^{-i\alpha\epsilon}$ with a constant α . Then*

$$\mathcal{L}_\psi[\psi', \bar{\psi}'] = \mathcal{L}_\psi[\psi, \bar{\psi}] - \alpha e (\partial_\mu \epsilon) J^\mu + \nabla_\mu (\dots), \quad (\text{A60})$$

so choosing $\alpha = 3c_V\eta$ ($= -\frac{3}{4}\eta$ in our normalization) cancels (A59) pointwise. Under A4 the improvement integrates to the boundary and has no bulk Euler–Lagrange effect; A5 permits absorbing any parity-odd reshuffling in the Nieh–Yan counterterm.

Sketch. Insert $\psi' = e^{i\alpha\epsilon}\psi$ into the kinetic term $e\bar{\psi}ie_A^\mu\gamma^A\partial_\mu\psi$ and use the Leibniz rule. The derivative acting on $e^{i\alpha\epsilon}$ generates the shift $-\alpha e(\partial_\mu\epsilon)J^\mu$ in (A60); the mass and spin-connection pieces are invariant under a *vector* (not axial) phase. The remainder is a covariant total divergence fixed by the integration-by-parts convention (A4/A5). $\square$

Combining (A59) with Prop. A5 and the coefficient assignment (A57) yields precisely the two-line summary in Section 8: no axial channel and a removable trace channel.

Path-integral measure (anomaly) check.

The field redefinition in Prop. A5 is a *vector* $U(1)$ rotation, hence the fermionic measure is invariant (Jacobian equal to one). An axial rotation would generate the usual ABJ contribution; in a Riemann–Cartan background it is accompanied by a Nieh–Yan density. We do not perform an axial rotation anywhere in this work.

Appendix E.3. Nieh–Yan as a Boundary Convention (A5)

The Nieh–Yan 4-form $NY \equiv d(e^A \wedge T_A)$ is an exact form. In our posture (A5), any explicit use of NY enters only as a *boundary convention* after applying the scalar projector Π_{PT} : it does not modify Euler–Lagrange equations in the bulk and is indistinguishable from a choice of improvement current. This applies equally to (i) the three-chain equivalence (where the PT -even CS/Nieh–Yan shadow differs from the DBI/CM routes by an improvement current) and (ii) the Dirac sector (where vector rephasing reshuffles boundary terms that can be absorbed into the chosen Nieh–Yan convention). No physical statement in Secs. 5–8 depends on a specific Nieh–Yan choice.

Appendix E.4. Scope and Caveats: Topology and Boundary Flux

Our conclusions rely on the posture A4: either compact PT -invariant domains with vanishing boundary flux or standard asymptotic fall-offs (FRW/flat) so that improvement currents integrate to zero. They also assume a *single-valued* spurion phase $\epsilon(x)$ so that the rephasing in Prop. A5 is a globally well-defined map $\psi \mapsto e^{i\alpha\epsilon}\psi$.

- **Nontrivial topology (not covered).** If ϵ is multi-valued or admits nontrivial holonomy (e.g., on manifolds with nontrivial π_1 or H^1), the map $e^{i\alpha\epsilon}$ may fail to be single-valued. In such cases the local cancellation in (A60) can leave a global remnant proportional to the winding; our bulk equivalences and removability statements are not asserted in that setting.
- **Nonvanishing boundary flux (not covered).** On domains where the PT -invariant boundary flux of the relevant improvement currents does not vanish, boundary terms may carry physical information (e.g., in explicitly finite boxes with prescribed inflow). Our null tests and cancellations are presented only under A4 fall-offs.

Appendix E.5. Bookkeeping Table and Quick References

For convenience we summarize the conventions and coefficients used in the Dirac sector:

Object / Convention	Value / Definition
Gamma matrices	$\{\gamma^A, \gamma^B\} = 2\eta^{AB}, \gamma_5 = i\gamma^0\gamma^1\gamma^2\gamma^3$
Hodge / Levi-Civita	$\epsilon^{0123} = +1, [\mathcal{PT}, *] = 0$ (A2)
Contorsion split	$K_{ABC} = \frac{1}{3}(\eta_{AC}T_B - \eta_{AB}T_C) - \frac{1}{6}\epsilon_{ABCD}S^D + q_{ABC}$
Dirac currents	$J^\mu = \bar{\psi}\gamma^\mu\psi, J_5^\mu = \bar{\psi}\gamma^\mu\gamma_5\psi$
Torsion→Dirac	$\Delta\mathcal{L}_\psi = e(\frac{3}{8}S_\mu J_5^\mu - \frac{1}{4}T_\mu J^\mu)$
C1 (pure trace)	$S^\mu = 0, q_{\lambda\mu\nu} = 0, T_\mu = 3\eta\partial_\mu\epsilon$
Vector rephasing	$\psi \rightarrow e^{i\alpha\epsilon}\psi$ with $\alpha = 3c_V\eta = -\frac{3}{4}\eta$
Outcome	Trace channel removed up to a total derivative (A4/A5)
Nieh–Yan	Boundary convention only (A5); no bulk Euler–Lagrange effect

Appendix E.6. Corollary: No LO Four-Fermion Contact from Torsion

At the order analyzed in this paper (quadratic in fields; at most one derivative per building block), torsion is algebraic and fixed by (C1), and the only linear Dirac–torsion channel that survives projection is removed by the vector redefinition above. Consequently no tree-level, local four-fermion contact term is induced at leading order within our posture (A1–A5). Any such effect would require (i) an axial channel ($S_\mu \neq 0$), (ii) higher-derivative completions beyond our closure basis, or (iii) loop corrections outside the present scope.

Summary of Appendix E. In our conventions the minimally coupled Dirac field interacts with torsion through $e(\frac{3}{8}S_\mu J_5^\mu - \frac{1}{4}T_\mu J^\mu)$. Under the pure-trace map (C1) the axial channel vanishes and the remaining trace channel is removed by a local vector rephasing $\psi \rightarrow e^{i(3c_V\eta)\epsilon}\psi$, up to a total derivative consistent with A4/A5. Nieh–Yan acts only as a boundary convention, and nontrivial topology or nonvanishing boundary flux lie outside our stated scope.

Appendix F. Operational Diagnostics (Extended)

This appendix collects a practitioner-oriented checklist and an expanded comparison table. Unless stated otherwise, all entries are framed for admissible patches with the boundary/topology posture A4, and for the quadratic order analyzed in the main text.

Appendix F.1. Full Decision Tree (D1–D9)

- **(D1) Parity channels.** Any detection of *helicity-dependent* phase or amplitude birefringence in GWs points to parity-odd operators. Our scalar-PT posture forbids such effects at LO (Thm. 1; Section 2.4).⁸
- **(D2) Axial/tensor torsion irreps.** Nonzero axial (S^μ) or traceless ($q_{\lambda\mu\nu}$) torsion signals are incompatible with C1. We predict the pure-trace map (Thm. 2); Figures 2, 4 provide diagnostics.
- **(D3) Three-route bulk collapse.** Reconstruct quadratic kernels from simulations/perturbation theory: the DBI, closed-metric, and CS^+ routes must share the same *bulk* piece $A_\star\sqrt{-g}\,\text{sgn}(\Sigma_\epsilon)\,I_T$ with $A_\star = \lambda^2/8$ (Section 5); residuals should enter only through improvements (Figure 5).
- **(D4) Boundary equivalence.** Boundary flux ratios $\mathcal{R}_{X/Y}[\partial\mathcal{D}]$ converge to 1 on growing domains (A4) (Section 5.2; Figure 6). Persistent deviations falsify C2.
- **(D5) Exact luminality by identity.** After eliminating TT–nonTT mixing with the full-rank 2×2 system (Section 6.2), the equal-coefficient identity $K - G = \partial_\mu(a^{-3}J_\Delta^\mu a^3)$ enforces $c_T=1$ without tuning (Section 7; Figure 9).
- **(D6) DoF count.** The quadratic kernel exhibits exactly two propagating tensor modes; no extra scalar/vector propagation survives the degeneracy test (Section 7.2; Figure 10). Any additional mode falsifies our posture at this order.

⁸ Chern–Simons modified gravity is the canonical benchmark for helicity-split propagation; see e.g. Jackiw–Pi; Alexander–Yunes.

Table A3. Operational disambiguation at quadratic order (expanded).

Diagnostic	This work	CS-mod. grav.	Horndeski /DHOST	Teleparallel / $f(T)$	EC/ ECSK-like
Parity-odd GW birefringence	No (D1)	Yes (generic)	No (often tuned)	Not diagnostic at LO	Model- dependent
Axial torsion S^μ at LO	Absent (C1)	Not diagnostic at LO	Not diagnostic at LO	Not diagnostic at LO	Often present
Traceless torsion $q_{\lambda\mu\nu}$ at LO	Absent (C1)	Not diagnostic at LO	Not diagnostic at LO	Possible in extensions	Possible
c_T at LO	= 1 by identity (D5)	Helicity-split	Tuned $\rightarrow$ 1	Often = 1	Model- dependent
Extra propagating DoF (quad.)	No (2 TT) (D6)	No new TT	Yes (scalar; generic)	Model- dependent	Model- dependent
Three-route bulk equality (C2)	Yes (D3)	Not applicable	Not applicable	Not applicable	Not applicable
Boundary flux ratio $\mathcal{R}_{X/Y}$	$\rightarrow$ 1 (D4)	Not applicable	Not applicable	Not applicable	Not applicable
NLO slope $\delta c_T^2 \propto k^2$	Yes (D7, D9-a)	Non-universal	Model- dependent	Model- dependent	Model- dependent
Dirac axial contact $S_\mu J_5^\mu$	Absent (D8)	Not diagnostic at LO	Not diagnostic at LO	Not diagnostic at LO	Present in general
Trace contact removable	Yes (D8)	Not diagnostic at LO	Not diagnostic at LO	Model- dependent	Not generically

- **(D7) NLO slope.** At next order, the leading PT -even dispersion predicts $\delta c_T^2(k) = b k^2 / \Lambda^2$ (Equation (38)): a log–log slope $\simeq 2$ in clean frequency windows is characteristic (Figure 11; Section 8.2).⁹
- **(D8) Fermion channel null.** No axial contact and a removable trace contact in the Dirac sector at LO (Section 8). Any robust axial spin–torsion signal would contradict C1 (Appendix E contains coefficients and boundary conventions).
- **(D9) Spurion-limit tests.** *Two complementary null checks* targeting residual spurion dynamics (Section 8.3):
 - **(D9-a) Quadratic scaling.** Fit $|\delta c_T^2(k)| \propto k^2$ on the admissible band ($k \ll \Lambda$) using Equation (39). A statistically significant failure ($\hat{n} \neq 2$ beyond systematics) indicates residual ϵ dynamics.
 - **(D9-b) Route equality.** Using the unified improvement representative (Appendix C), test $\Delta A_\star(k) \equiv A_\star^{(\text{rank} - \text{onedeterminantroute})} - A_\star^{(\text{CM})} \rightarrow 0$ and boundary flux-ratio unity $\mathcal{R}_{\text{rank} - \text{onedeterminantroute}/\text{CM}} \rightarrow 1$ on growing domains (A4). Significant residuals falsify the spurion limit.

Appendix F.2. Expanded Disambiguation Table

Entries are *generic at LO* on admissible backgrounds with A4 fall-offs; tuned subclasses and specific parameter choices may alter individual cells. See the notes in Section F.3 for caveats and representative citations.

Appendix F.3. Notes, Caveats, and Representative Anchors

Parity-odd lines.

Chern–Simons modified gravity (Jackiw–Pi; Alexander–Yunes) generically induces helicity-dependent GW propagation; details depend on the choice of scalar field and coupling. Our scalar- PT projector eliminates parity-odd *scalars* at LO; see Thm. 1. Representative anchors: [25,26].

⁹ Caveat: finite-window fits may be biased by instrumental systematics or astrophysical priors; interpret bounds conservatively.

Horndeski/DHOST.

Post-GW170817 constraints force $c_T \rightarrow 1$ by parameter tuning or by restricting to subclasses; however, no equal-coefficient *identity* of the type $K-G = \partial_\mu(\dots)$ is generally present, and an extra scalar DoF is typical. See the multimessenger bounds and reviews [31–36].

Teleparallel/ $f(T)$.

Many models yield $c_T = 1$ at LO; torsion irreps and matter couplings are model-dependent, and our route-equality/flux-ratio diagnostics (C2) do not directly apply. We therefore treat teleparallel cases as outside the present “three-route” posture.

EC/ECSK-like.

With propagating torsion or nonminimal fermion couplings, axial/traceless torsion irreps can be present and Dirac axial contacts typically survive. Our C1 pure-trace map excludes these at the analyzed order; see classic EC/MAG treatments [4–6,8].

NLO dispersion and spurion tests.

The $\delta c_T^2 \propto k^2$ slope follows from the projected, PT -even closure and the normalization of $\mathcal{T}$ (Section 5); D9-a implements this as a band-level null. D9-b leverages the three-route equality (Appendix C) via $\Delta A_\star$ and flux ratios; in admissible domains (A4) both must vanish within errors.

Boundary posture.

All boundary-sensitive statements assume A4 (PT-invariant compact boundaries or standard AF/FRW fall-offs) and A5 (Nieh–Yan as boundary counterterm). Departures from these postures may alter flux diagnostics and improvement accounting. Operational anchors: covariant phase space and surface charges [13–15], and the Holst/Nieh–Yan parity bookkeeping [9–12].

Representative anchors (one-line map).

Metric-affine/Einstein–Cartan: [4,6,8]. Holst/Nieh–Yan: [9–12]. Covariant phase space & boundary charges: [13–15]. CS-modified gravity: [25,26]. GW170817 speed bounds & implications: [27–34].

Appendix G. EFT Origins of the Spurion Field

Scope. This appendix provides an effective-field-theory (EFT) completion that *explains* the spurion posture adopted in the main text and ties it to *projective symmetry*. At two derivatives and within A1–A5 (Section 2), the compensator $\epsilon(x)$ appears only through its gradient and *only* in the projectively invariant combination

$$\mathcal{T}_\mu \equiv T_\mu - \partial_\mu \epsilon, \quad \Gamma^\alpha_{\mu\nu} \rightarrow \Gamma^\alpha_{\mu\nu} + \delta^\alpha_\mu \xi_\nu : \mathcal{T}_\mu \text{ invariant.} \quad (\text{A61})$$

We give a minimal Stueckelberg action that reproduces the C1 map, explain two equivalent implementations of the *spurion limit*, and record complementary EFT viewpoints in which the gradient-only appearance of ϵ is manifest.

Appendix G.1. G.1 Stueckelberg Completion and the Spurion Limit

A projectively invariant, two-derivative completion is

$$S_{\text{Stk}} = \int d^4x \sqrt{-g} \left[\frac{M_P^2}{2} R(e, \omega) - \frac{m_T^2}{2} \mathcal{T}_\mu \mathcal{T}^\mu - \frac{f_\epsilon^2}{2} (\partial\epsilon)^2 \right], \quad \mathcal{T}_\mu \equiv T_\mu - \partial_\mu \epsilon, \quad (\text{A62})$$

with $m_T > 0$ and $f_\epsilon \geq 0$ real. The masslike term for $\mathcal{T}_\mu$ restores projective invariance via the Stueckelberg compensator ϵ ; the spectator kinetic term respects the shift symmetry $\epsilon \mapsto \epsilon + c$ and keeps only $\partial\epsilon$ at this order.

Palatini variation and the C1 map.

Writing $\omega = \dot{\omega} + K(T)$ and using the standard algebraic split of $R(e, \omega)$ into $R(\dot{\omega})$ plus quadratic torsion (total derivatives dropped under A4), variation w.r.t. ω yields purely *algebraic* equations in the irreps $\{T_\mu, S^\mu, q_{\lambda\mu\nu}\}$. With (A62),

$$\delta_\omega S_{\text{Stk}} \Rightarrow S^\mu = 0, \quad q_{\lambda\mu\nu} = 0, \quad m_T^2 (T_\mu - \partial_\mu \epsilon) + \frac{4\alpha_1}{3} T_\mu = 0, \quad (\text{A63})$$

where α_1 is the coefficient multiplying the T^2 piece coming from $R(e, \omega)$ (fixed by conventions and already used in Section 4.1). For $m_T^2 > 0$ this algebraic system has the unique solution

$$\boxed{T_\mu = \partial_\mu \epsilon, \quad S^\mu = 0, \quad q_{\lambda\mu\nu} = 0,} \quad (\text{A64})$$

which is precisely the C1 *pure-trace alignment* (Section 4) written in covector form and implies $\mathcal{T}_\mu = 0$ on-shell.¹⁰

Two equivalent implementations of the *spurion limit*.

We isolate the low-energy regime where observables reduce to those used in the main text:

$$\boxed{\begin{array}{l} \text{(i) Hard penalty: } m_T \rightarrow \infty \Rightarrow \mathcal{T}_\mu = 0, \\ \text{(ii) Lagrange current: } \mathcal{L}_\Lambda = \sqrt{-g} \Lambda^\mu \mathcal{T}_\mu. \end{array}} \quad (\text{A65})$$

In case (i) integrating out ω (equivalently T) produces a functional delta $\delta[\mathcal{T}]$; in case (ii) the constraint is exact already at the classical level. In both realizations, keeping f_ϵ^2 finite leaves a standard $(\partial\epsilon)^2$ spectator; the *spurion posture* used in the main text corresponds to taking $f_\epsilon^2 \rightarrow 0^+$ (or treating ϵ as nondynamical) *within the two-derivative truncation*, so that observables depend only on $\partial\epsilon$ through the invariant $\mathcal{T}_\mu$ (and hence through Σ_ϵ once C1 is enforced).

Low-energy effective action (power counting).

At energies $E \ll m_T$ one obtains, after eliminating ω ,

$$S_{\text{eff}}^{(\text{LO})} = \int d^4x \sqrt{-g} \left[\frac{M_P^2}{2} R(\dot{\omega}) + \alpha_1 I_T + \alpha_2 \Sigma_\epsilon \right] + \mathcal{O}\left(\frac{\partial^4}{m_T^2}\right), \quad I_T \equiv -\frac{1}{4} \Pi_{\text{PT}}[T^A{}_{BC} T_A{}^{BC}], \quad (\text{A66})$$

and *after* enforcing C1 (Section 4), $I_T = -6\eta^2 \Sigma_\epsilon$, so that the LO bulk reduces to a single invariant line in the $\{I_T, \Sigma_\epsilon\}$ plane (Section 3.3).

Appendix G.2. Complementary EFT Pictures (Same Gradient-Only Physics)

Two complementary constructions emphasize that *only* $\partial\epsilon$ can enter observables at this order.

(i) Torsional–axion picture.

Augment (A62) by a topological coupling

$$S_{\text{t-ax}} = \int d^4x \sqrt{-g} \frac{\alpha}{4} \epsilon \text{NY}, \quad \text{NY} \equiv d(e^A \wedge T_A), \quad (\text{A67})$$

which is exact. Under A5, NY is a *boundary convention*; integrating by parts and using (A64) gives bulk terms proportional to $\partial_\mu \epsilon$ contracted with improvement currents, hence no new bulk Euler–Lagrange content in the PT -even sector. The low-energy reduction again depends only on Σ_ϵ (plus boundary assignments absorbed into A5).

¹⁰ Equivalently, one may add the quadratic invariant I_T with a positive coefficient and obtain the same irrep alignment. The projector statements of Section 2 ensure all scalar contractions are taken in the PT -even, real sector.

(ii) Axion electrodynamics-like geometry.

The situation mirrors $\theta F\tilde{F}$ in gauge theory: the integrated density is topological, and only $\partial_\mu\theta$ is physically measurable (through Chern–Simons currents) at two derivatives. Replacing $(\theta, F\tilde{F})$ by (ϵ, NY) and using the scalar PT projector (A2) reproduces the same *gradient-only* observability for ϵ .

Appendix G.3. UV Hints and Emergence

While our analysis is agnostic about ultraviolet completions, two standard paths naturally generate a field with the required properties:

1. **Torsionful connections in string-inspired setups.** In heterotic/type II supergravities the NS–NS three-form $H = dB$ enters via the *torsionful* spin connection $\omega_\pm = \hat{\omega} \pm \frac{1}{2} H$. In slowly varying, parity-even sectors and upon dimensional reduction, gradients of the axion-like fields dual to B act as effective *trace* torsion, yielding ∂_μ of a scalar in the infrared and a projectively invariant combination akin to $\mathcal{T}_\mu$.
2. **'t Hooft naturalness from symmetry.** The projective symmetry together with the shift symmetry $\epsilon \rightarrow \epsilon + c$ protects the *gradient-only* appearance of ϵ and suppresses f_ϵ^2 technically: taking $f_\epsilon^2 \rightarrow 0$ *enhances* symmetry (exact spurion posture).

These UV *hints* are illustrative only; no UV assumption is needed for the main results.

Appendix G.4. Remarks on Matter Couplings and Boundary

Dirac sector at LO (consistency).

Adding minimally coupled fermions leaves the conclusions unchanged at this order: C1 implies $S^\mu=0$, $q_{\lambda\mu\nu}=0$ and $T_\mu \parallel \partial_\mu\epsilon$; the surviving $T_\mu J^\mu$ contact is removed by a *vector* rephasing $\psi \rightarrow e^{ia\epsilon}\psi$ (Appendix E) up to improvements consistent with A4/A5. Thus no axial contact and no irreducible trace contact survive at LO.

Boundary/topology posture.

All reductions above use A4 (compact PT -invariant boundaries or standard AF/FRW fall-offs) and A5 (Nieh–Yan as a boundary counterterm). Nontrivial holonomies of ϵ or prescribed boundary inflows can carry global information not analyzed here.

Summary and Cross-Reference

Appendix G in one line. A minimal, projectively invariant Stueckelberg completion $S_{\text{Stk}} = \int \sqrt{-g} \left[\frac{M_P^2}{2} R - \frac{m_T^2}{2} (T - \partial\epsilon)^2 - \frac{f_\epsilon^2}{2} (\partial\epsilon)^2 \right]$ algebraically enforces the pure-trace alignment $T_\mu = \partial_\mu\epsilon$ (C1) and $S^\mu = q_{\lambda\mu\nu} = 0$. Taking the *spurion limit* (hard penalty $m_T \rightarrow \infty$ or a Lagrange current for $\mathcal{T}_\mu$, with $f_\epsilon^2 \rightarrow 0$ at two derivatives) yields precisely the observable sector used in the main text: PT -even scalars depend only on $\mathcal{T}_\mu$ (hence on Σ_ϵ once C1 is in force), and the LO bulk reduces to $\alpha_1 I_T + \alpha_2 \Sigma_\epsilon$ with $I_T = -6\eta^2 \Sigma_\epsilon$. Complementary EFT pictures (torsional–axion; axion–ED analogy) lead to the same gradient-only dependence.

Hook to Sec. 9. The statement used in Section 9—“Treating ϵ as non-dynamical is the low-energy limit of a Stueckelberg completion in which $m_T \rightarrow \infty$ freezes $\mathcal{T}_\mu$; complementary EFT constructions reduce to the same gradient-only dependence”—is precisely the content of Eqs. (A62)–(A65).

References

1. A. Palatini, “Deduzione invariantiva delle equazioni gravitazionali dal principio di Hamilton,” *Rend. Circ. Mat. Palermo* **43**, 203–212 (1919).
2. T. W. B. Kibble, “Lorentz invariance and the gravitational field,” *J. Math. Phys.* **2**, 212–221 (1961).
3. D. W. Sciama, “On the analogy between charge and spin in general relativity,” in *Recent Developments in General Relativity* (Pergamon, 1962), pp. 415–439.

4. F. W. Hehl, P. von der Heyde, G. D. Kerlick, and J. M. Nester, "General Relativity with Spin and Torsion: Foundations and Prospects," *Rev. Mod. Phys.* **48**, 393–416 (1976).
5. F. W. Hehl and B. K. Datta, "Nonlinear spinor equation and asymmetric connection in general relativity," *J. Math. Phys.* **12**, 1334–1339 (1971).
6. F. W. Hehl, J. D. McCrea, E. W. Mielke, and Y. Ne'eman, "Metric-affine gauge theory of gravity: Field equations, Noether identities, world spinors, and breaking of dilaton invariance," *Phys. Rept.* **258**, 1–171 (1995).
7. R. T. Hammond, "Torsion gravity," *Rep. Prog. Phys.* **65**, 599–649 (2002).
8. I. L. Shapiro, "Physical aspects of the space–time torsion," *Phys. Rept.* **357**, 113–213 (2002).
9. H. T. Nieh and M. L. Yan, "An identity in Riemann–Cartan geometry," *J. Math. Phys.* **23**, 373–374 (1982).
10. O. Chandia and J. Zanelli, "Topological invariants, instantons and the chiral anomaly on spaces with torsion," *Phys. Rev. D* **55**, 7580–7585 (1997).
11. S. Holst, "Barbero's Hamiltonian derived from a generalized Hilbert–Palatini action," *Phys. Rev. D* **53**, 5966–5969 (1996).
12. S. Mercuri, "Fermions in the Ashtekar–Barbero connection formalism: The Nieh–Yan invariant as a source of the Immirzi parameter," *Phys. Rev. D* **73**, 084016 (2006).
13. R. M. Wald, *General Relativity* (University of Chicago Press, 1984).
14. V. Iyer and R. M. Wald, "Some properties of Noether charge and a proposal for dynamical black hole entropy," *Phys. Rev. D* **50**, 846–864 (1994).
15. T. Regge and C. Teitelboim, "Role of surface integrals in the Hamiltonian formulation of general relativity," *Annals Phys.* **88**, 286–318 (1974).
16. Yu. N. Obukhov, "Poincaré gauge gravity: Selected topics," *Int. J. Geom. Methods Mod. Phys.* **3**, 95–138 (2006).
17. Y. N. Obukhov and F. W. Hehl, "Rotation, acceleration, and gravity in the framework of classical electrodynamics," *Phys. Lett. A* **372**, 3946–3952 (2008); see also F. W. Hehl and Y. N. Obukhov, *Foundations of Classical Electrodynamics* (Birkhäuser, 2003), App. B.
18. T. P. Sotiriou and V. Faraoni, " $f(R)$ theories of gravity," *Rev. Mod. Phys.* **82**, 451–497 (2010).
19. G. J. Olmo, "Palatini Approach to Modified Gravity: $f(R)$ Theories and Beyond," *Int. J. Mod. Phys. D* **20**, 413–462 (2011).
20. E. E. Flanagan, "Palatini form of $1/R$ gravity," *Phys. Rev. Lett.* **92**, 071101 (2004).
21. E. Barausse, T. P. Sotiriou, and J. C. Miller, "Curvature singularities in Palatini $f(R)$ gravity," *Phys. Rev. D* **77**, 104035 (2008).
22. J. D. Bekenstein, "The Relation between physical and gravitational geometry," *Phys. Rev. D* **48**, 3641–3647 (1993).
23. S. Deser and G. W. Gibbons, "Born–Infeld–Einstein actions?," *Class. Quantum Grav.* **15**, L35–L39 (1998).
24. M. Bañados and P. G. Ferreira, "Eddington's theory of gravity and its progeny," *Phys. Rev. Lett.* **105**, 011101 (2010).
25. R. Jackiw and S.-Y. Pi, "Chern–Simons modification of general relativity," *Phys. Rev. D* **68**, 104012 (2003).
26. S. Alexander and N. Yunes, "Chern–Simons Modified Gravity," *Phys. Rept.* **480**, 1–55 (2009).
27. B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), "GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral," *Phys. Rev. Lett.* **119**, 161101 (2017).
28. B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), "Multi-messenger Observations of a Binary Neutron Star Merger," *Astrophys. J. Lett.* **848**, L12 (2017).
29. A. Goldstein *et al.*, "An Ordinary Short Gamma-Ray Burst with Extraordinary Implications: Fermi-GBM Detection of GRB 170817A," *Astrophys. J. Lett.* **848**, L14 (2017).
30. V. Savchenko *et al.*, "INTEGRAL Detection of the First Prompt Gamma-Ray Signal Coincident with the Gravitational-wave Event GW170817," *Astrophys. J. Lett.* **848**, L15 (2017).
31. P. Creminelli and F. Vernizzi, "Dark Energy after GW170817 and GRB170817A," *Phys. Rev. Lett.* **119**, 251302 (2017).
32. T. Baker, E. Bellini, P. G. Ferreira, M. Lagos, J. Noller, and I. Sawicki, "Strong constraints on cosmological gravity from GW170817 and GRB170817A," *Phys. Rev. Lett.* **119**, 251301 (2017).
33. J. M. Ezquiaga and M. Zumalacárregui, "Dark Energy After GW170817: Dead Ends and the Road Ahead," *Phys. Rev. Lett.* **119**, 251304 (2017).
34. J. Sakstein and B. Jain, "Implications of the Neutron Star Merger GW170817 for Cosmological Scalar–Tensor Theories," *Phys. Rev. Lett.* **119**, 251303 (2017).
35. R. Kase and S. Tsujikawa, "Dark energy in Horndeski theories after GW170817: A review," *Int. J. Mod. Phys. D* **28**, 1942005 (2019).

36. D. Langlois, "Dark Energy and Modified Gravity in Degenerate Higher-Order Scalar–Tensor (DHOST) theories," *Int. J. Mod. Phys. D* **28**, 1942006 (2019).
37. Yu. N. Obukhov, "Gravitational waves in Poincaré gauge gravity theory," *Phys. Rev. D* **95**, 084028 (2017).
38. E. Elizalde, S. D. Odintsov, and V. V. Obukhov, "Gravitational waves in Einstein–Cartan theory," *Phys. Dark Univ.* **41**, 101256 (2023).
39. N. V. Agazie *et al.* (NANOGrav Collaboration), "The NANOGrav 15-year Data Set: Evidence for a Gravitational-Wave Background," *Astrophys. J. Lett.* **951**, L8 (2023).
40. J. Antoniadis *et al.* (EPTA Collaboration and InPTA Collaboration), "The second data release from the European Pulsar Timing Array: Search for signals from new physics," *Astron. Astrophys.* **678**, A50 (2023).
41. D. J. Reardon *et al.* (PPTA Collaboration), "Search for an Isotropic Gravitational-Wave Background with the Parkes Pulsar Timing Array," *Astrophys. J. Lett.* **951**, L7 (2023).
42. F. A. A. Liu *et al.* (CPTA Collaboration), "Searching for the nano-Hertz stochastic gravitational-wave background with the Chinese Pulsar Timing Array Data Release I," *Res. Astron. Astrophys.* **23**, 075024 (2023).
43. R. Arnowitt, S. Deser, and C. W. Misner, "The Dynamics of General Relativity," in *Gravitation: An Introduction to Current Research*, ed. L. Witten (Wiley, 1962), pp. 227–265; reprinted in *Gen. Rel. Grav.* **40**, 1997–2027 (2008).

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.