symmetry

IMPACT FACTOR 2.2

CITESCORE 5.3

Article

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

Chien-Chih Chen

Special Issue
Symmetry, Topology and Geometry in Physics

Edited by
Prof. Dr. Gabino Estevez-Delgado

MDPI

https://doi.org/10.3390/sym18010170

symmetry

MDPI

Article

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

Chien-Chih Chen

Chunghwa Telecom Laboratories, Information & Communications Security Laboratory, Chunghwa Telecom Co., Ltd., Taoyuan City 326, Taiwan; 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 and Scope. *Positioning and scope (read first).* What this paper is: *a metric-compatible first-order (Palatini) variational formulation on a Riemann–Cartan spacetime with torsion, equipped with projective invariance and a scalar PT projection, proving a symmetry-selected luminality mechanism in the quadratic tensor sector.* What this paper is not: *(i) not teleparallel gravity (we do not impose vanishing curvature); (ii) not a comprehensive modified gravity survey (we cite only minimal landmarks needed to locate the contribution); (iii) not a full derivation/solution of background cosmology or parameter fits. Our aim is a specification-level statement: under explicit assumptions (A1–A6), $c_T = 1$ emerges structurally, and the leading falsifiable deviation is pushed to NLO in a controlled, symmetry-consistent operator basis.*

The dawn of multimessenger astronomy, heralded by the binary-neutron-star event GW170817 and its electromagnetic counterparts, has transformed the empirical status of gravity: the near-simultaneous arrival of gravitational waves and gamma rays constrains the tensor propagation speed to be luminal to extraordinary precision,

Academic Editor: Stefano Profumo

Received: 20 December 2025
Revised: 9 January 2026
Accepted: 13 January 2026
Published: 16 January 2026

$|c_T/c - 1| \lesssim 10^{-15}$ [1–3]. This single datum acts as a stringent filter across a wide landscape of modified gravity models, motivating a sharp question: *can luminality follow from symmetry and consistency, rather than from parameter tuning?* A parameter-tuned survival of luminality is often viewed as theoretically fragile; a symmetry-selected mechanism, by contrast, would be structural and radiatively meaningful. For reproducibility, all numerical checks and figure generation were performed using Python (version 3.11.6) with standard scientific libraries; the public repository was accessed on 19 December 2025.

Related work map (minimal, not a survey).

The GW170817 constraint rapidly eliminated broad classes of scalar–tensor theories unless specific operators are absent or degenerately arranged [4–6]. Representative EFT overviews of viable post-GW170817 parameterizations can be found in [7–10]. On the geometric side, first-order (Palatini) and metric-affine/Riemann–Cartan formulations with torsion have a long history [11–14]. The spin–torsion coupling in the Einstein–Cartan/Sciama–Kibble framework was introduced in [15,16]. For an overview of torsion gravity and related formulations, see [17]. For clarity, we emphasize again that our setting is *not* teleparallel gravity (which enforces zero curvature and attributes gravity to torsion alone); we cite standard teleparallel reviews solely to prevent misclassification [18,19]. No result in this paper relies on teleparallel assumptions.

Framework and posture.

We work in a metric-compatible, first-order (Palatini) variational formulation where the tetrad (or metric) and spin connection are treated as independent variables on a Riemann–Cartan spacetime (torsion allowed). Within this setting, we impose two symmetry/selection principles that define the "observable" scalar sector and control the tensor dynamics:

1. **Scalar PT projection.** All observable scalar densities are projected into a PT-even (parity–time symmetric) subspace. Operationally, this acts as a selection rule on admissible scalar operators in the action, ensuring reality of the action and removing parity-odd scalar contributions (e.g., potential birefringent effects) at leading order within the observable sector.
2. **Projective invariance with a compensator.** Projective symmetry is implemented with a non-dynamical compensator (spurion) $\epsilon(x)$, such that observables depend only on the projectively invariant trace combination

$$\mathcal{T}_\mu \equiv T_\mu - \partial_\mu \epsilon \,, \tag{1}$$

thereby avoiding additional propagating degrees of freedom that commonly arise in torsionful settings. Throughout the main analysis, we adopt the spurion limit (A6): ϵ is treated as a fixed background and is not varied.

These principles, together with six explicit assumptions (A1–A6) collected in Section 2, define the posture in which the quadratic tensor sector becomes rigid enough to admit an algebraic luminality theorem. Detailed proofs regarding the mathematical completeness of the operator basis and the local no-go theorems underpinning this posture are provided in the companion technical note [20].

Known ingredients vs. new contributions (audit-ready).

To make originality and logical dependencies explicit, we separate standard tools from new results.

- **Known/standard ingredients:** first-order variational methods in gravity [11]; Riemann–Cartan torsion decomposition and torsionful gravity foundations [12–14]; post-GW170817 luminality constraints and EFT framing [4–8].
- **This paper's new contributions (quadratic and falsifiable):**
  - **C1** a **uniqueness map** enforcing a pure-trace torsion alignment $T_\mu \parallel \partial_\mu \epsilon$ under the posture (A1–A6) (see [20] for the proof of uniqueness among all admissible 2-derivative operators);
  - **C2** a **three-route bulk equivalence** between distinct constructions, holding modulo boundary terms;
  - **C3** a **coefficient locking mechanism** yielding $K = G$ in the quadratic tensor action and thus $c_T = 1$ **by construction** (without parameter tuning);
  - **NLO** a **symmetry-consistent operator basis at four-derivative order** that isolates the *unique* source of luminality breaking, leading to a sharp dispersion target testable with multi-band GW data.

Quadratic luminality: Symmetry-selected locking.

Within the above posture, the logical chain is summarized in Figure 1. The PT selection and projective structure constrain torsion to a unique pure-trace form (C1), enable bulk equivalences between three construction routes up to $\nabla_\mu J^\mu$ (C2), and culminate in a locking condition that enforces $K = G$ and hence exact luminality $c_T = 1$ (C3). The key point is structural: luminality arises from consistency in the quadratic tensor sector, not from tuning of free parameters.

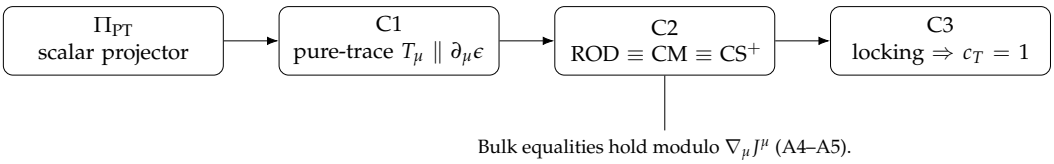

Figure 1. Roadmap of the quadratic analysis (bulk-only). The scalar *PT* projector (Theorem 1) defines the allowed observable sector. Projective invariance is realized via $\mathcal{T}_\mu = T_\mu - \partial_\mu \epsilon$. (C1) The posture enforces a pure-trace torsion aligned with the spurion gradient (Section 4, Equation (16)). (C2) It establishes the bulk equivalence of three construction routes up to boundary terms (Section 5, Equation (28)). (C3) A locking condition then yields $K = G$ and thus $c_T = 1$ by construction (Section 6, Equation (39)).

Falsifiable NLO signature and uniqueness at four derivatives.

The leading deviation from luminality is pushed to next-to-leading order (NLO), taking the universal EFT form

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

where k is the wavenumber and Λ is the effective scale of new physics. Crucially, in the symmetry-restricted four-derivative operator basis, *only one operator direction* generates a genuine $K \neq G$ mismatch in the transverse-traceless reduction: the Weyl-squared term. All other admissible four-derivative scalars either renormalize K and G in lockstep (preserving luminality) or reduce, under the admissible equivalences (A4–A5), to boundary/improvement structures. This "single breaking direction" makes the NLO prediction technically stable and model-selective: any radiative destabilization of $c_T = 1$ must enter through the same k^2-suppressed channel. We provide an operator-by-operator TT reduction to audit this claim (see Appendix H for the complete accounting). A detection (or robust

non-detection) of the specific k^2 scaling (2) across PTA, LISA, and ground-based bands would therefore provide a sharp, falsifiable test of the framework.

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

Positioning and Scope. *This section specifies the* posture *of the paper as a precise working specification. We fix the variational setting, define the symmetry-selected observable sector, and state the explicit assumptions (A1–A6) under which all later claims are to be read. In particular, we (i) work in a metric-compatible first-order (Palatini) variational formulation on a Riemann–Cartan spacetime with torsion, (ii) impose projective invariance in observables via a Stueckelberg spurion, and (iii) apply a scalar PT projection to observable scalar densities. We do* not *assume teleparallelism (no zero-curvature constraint), and we do not attempt a full background solution or parameter fitting.*

2.1. Terminology Note: Variational Principle vs. Geometry

We clarify terminology to prevent a common misreading. "Palatini" refers to a first-order variational principle in which the coframe/tetrad e^A (or metric) and the connection ω^{AB} are treated as independent variables [11–13]. The underlying spacetime geometry in this paper is Riemann–Cartan: curvature is generally nonzero and torsion is allowed. We furthermore restrict to metric-compatible connections, so $\nabla_\mu g_{\alpha\beta} = 0$.

Projective invariance is a standard redundancy of Palatini/metric-affine descriptions: a shift of the connection by a pure-trace piece does not affect certain curvature invariants and can be treated as a gauge-like redundancy in the observable sector. Finally, we emphasize that our setting is not teleparallel gravity: teleparallel models impose vanishing curvature and attribute gravity entirely to torsion. We do not impose such a constraint; teleparallel reviews are cited only to preclude misclassification [18,19].

2.2. Geometric Setting and Conventions

We work on an oriented $(3+1)$-dimensional Lorentzian manifold. In the first-order formalism, the basic variables are a coframe e^A and a metric-compatible spin connection $\omega^{AB} = -\omega^{BA}$. The torsion and curvature two-forms are

$$T^A \equiv de^A + \omega^A{}_B \wedge e^B, \qquad R^{AB} \equiv d\omega^{AB} + \omega^A{}_C \wedge \omega^{CB}. \tag{3}$$

Component conventions follow standard Riemann–Cartan notation [12–14].

Throughout, $\epsilon^{\mu\nu\rho\sigma}$ denotes the Levi–Civita tensor density compatible with the chosen orientation.

2.3. Principle 1: Projective Symmetry and the Stueckelberg Spurion

Our first guiding principle is projective invariance in observables. At the level of an affine connection, a projective transformation acts as

$$\Gamma^\alpha{}_{\mu\nu} \;\to\; \Gamma^\alpha{}_{\mu\nu} + \delta^\alpha_\mu \xi_\nu, \qquad T_\mu \equiv T^\alpha{}_{\mu\alpha} \;\to\; T_\mu + (d-1)\,\xi_\mu, \tag{4}$$

where $d = 4$ and ξ_μ is an arbitrary one-form. To ensure that physical observables are insensitive to this redundancy, we introduce a Stueckelberg compensator (spurion) field $\epsilon(x)$ with the compensating transformation $\partial_\mu\epsilon \to \partial_\mu\epsilon + (d-1)\,\xi_\mu$. This packages all projectively invariant dependence of the torsion trace into the combination

$$\boxed{\mathcal{T}_\mu \equiv T_\mu - \partial_\mu\epsilon} \quad \text{(projectively invariant).} \tag{5}$$

Spurion Limit (A6) and "What We Do/Do Not Vary"

Within the two-derivative truncation adopted in this paper, we treat ϵ as a non-dynamical background (spurion): it is held fixed and is not varied in the action principle at the order considered. This is a deliberate EFT truncation choice that preserves projective invariance while keeping the quadratic analysis auditable. If ϵ acquires small dynamics (e.g., a kinetic term $Z_\epsilon(\partial\epsilon)^2$ or higher-derivative operators), it introduces an additional scalar channel that can couple to the tensor sector only through symmetry-allowed higher-order operators; we return to this point when discussing NLO signatures and null tests in Section 8.

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

Positioning and Scope. *This subsection is interpretive and optional. It provides one UV narrative in which a phase-like spurion ϵ and a scalar-level PT selection rule arise naturally. No proof in this paper depends on this narrative; all structural results (C1–C3) follow solely from the posture A1–A6 and standard variational identities.*

A possible UV reading is that ϵ behaves as a phase/clock-like geometric spurion and that the scalar PT projection implements a minimal reality filter on the observable IR sector. Such a viewpoint can guide model-building beyond the conservative truncation adopted here, but it does not enter any derivation of the quadratic luminality mechanism.

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

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

Why PT on Scalar Densities?

Imposing PT as a selection rule on scalar densities is the minimal intervention compatible with the first-order variational framework. It does not require constraining all microscopic components of (e, ω), yet it ensures that any scalar density entering the action belongs to the PT-even observable sector once projected.

2.6. 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. *PT-invariant domain/measure.* For any scalar density X, the integration domain and measure obey $\int X = \int \mathcal{PT}[X]$.

A2. *Orientation and Hodge dual.* The combined PT transformation preserves spacetime orientation and commutes with the Hodge star, $[\mathcal{PT}, *] = 0$.

A3. *Projection–variation commutation (scalar densities).* For scalar densities $\mathcal{O}$, variation and projection commute on the admissible domain: $\delta\,\Pi_{\rm PT}[\mathcal{O}] = \Pi_{\rm PT}[\delta\mathcal{O}]$ (Lemma 1).

A4. *Admissible boundaries/fall-offs.* Improvement currents contribute only boundary terms and do not modify the symplectic form (e.g., compact domains or standard asymptotic fall-offs).

A5. *Nieh–Yan as a boundary term.* The Nieh–Yan 4-form $\mathrm{NY} \equiv d(e^A \wedge T_A)$ contributes only through boundary terms [21,22].

A6. *Two-derivative truncation and spurion limit.* We restrict to quadratic order with at most two derivatives in the tensor sector. The compensator ϵ is treated as a fixed background (spurion) at this order and is not varied.

Admissible Domains and Caveats

Outside A4 (e.g., configurations with torsion defects or nontrivial global spurion holonomy), boundary fluxes can obstruct bulk equivalences. All equivalence statements in later sections are therefore to be read as bulk statements modulo total derivatives on admissible domains.

2.7. Formal Properties of the PT Projector

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

$$\Pi_{\rm PT}[\mathcal{O}] \equiv \frac{1}{2}(\mathcal{O} + \mathcal{PT}[\mathcal{O}]), \qquad \Pi_{\rm PT}^2 = \Pi_{\rm PT}. \tag{6}$$

Assumption A1 ensures that this projection defines a consistent restriction at the level of integrated observables. We treat time reversal T as anti-unitary: it includes complex conjugation acting on scalar coefficients.

Lemma 1 (Projection–variation commutation (A3))**.** *Under A1 and A4, with the spurion ϵ held fixed, one has $\delta\,\Pi_{\rm PT}[\mathcal{O}] = \Pi_{\rm PT}[\delta\mathcal{O}]$ for any scalar density $\mathcal{O}$ under variations of (e, ω).*

Proof. Linearity yields $\delta\,\Pi_{\rm PT}[\mathcal{O}] = \frac{1}{2}(\delta\mathcal{O} + \delta\,\mathcal{PT}[\mathcal{O}])$. On admissible domains (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. Hence, $\delta\,\mathcal{PT}[\mathcal{O}] = \mathcal{PT}[\delta\mathcal{O}]$, proving the claim. □

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

Proof. Immediate from (6). □

2.8. Discrete-Action Conventions in the Projected Sector

For reference, Table 1 summarizes a consistent choice of P and T actions on representative objects. This table is used only to enforce scalar density selection rules; it does not impose microscopic constraints on all components of (e, ω).

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

Object	P	T	Note
${e^A}_\mu$ (vierbein)	$+$	$+$	
${\omega^{AB}}_\mu$ (connection)	$+$	$-$	T acts anti-unitarily
${T^A}_{\mu\nu}$ (torsion)	$+$	$-$	
${R^{AB}}_{\mu\nu}$ (curvature)	$+$	$+$	
ϵ (spurion)	$+$	$-$	chosen to keep scalar sector consistent
$\partial_\mu\epsilon$	$-$	$+$	vector-like under P
$\epsilon^{\mu\nu\rho\sigma}$ (density)	$-$	$-$	PT-even

2.9. Variational Outputs at the Order Considered

We record the variational outputs needed for later sections, while keeping the scope explicit. At the order considered, we vary only with respect to the geometric variables (e, ω), with ϵ held fixed (spurion limit).

Variational outputs (posture A1–A6).

- **Variation w.r.t. ω (connection):** yields algebraic torsion equations at two-derivative order, which under A1–A6 enforce the pure-trace alignment used in Result C1 (Section 4).
- **Variation w.r.t. e (coframe/metric):** yields the metric field equations. We do not attempt to solve background cosmologies in this paper; the quadratic tensor speed proof depends only on the reduced quadratic action and the admissible-domain identities (A4–A5).
- **Variation w.r.t. ϵ:** in the spurion limit, ϵ is not varied. If ϵ is promoted to a dynamical field in an extended EFT, its equation takes a conserved-current/constraint form (e.g., $\nabla_\mu J^\mu_\epsilon = 0$ for an appropriate current J^μ_ϵ), and its couplings to the tensor sector enter only through symmetry-allowed higher-order operators; this lies beyond the present scope.

2.10. Summary of Key Observable-Sector Ingredients

We consolidate the spurion-built and projected quantities that appear repeatedly in later sections:

Key ingredients for the observable sector

- *Projectively invariant trace:* $\mathcal{T}_\mu \equiv T_\mu - \partial_\mu \epsilon$.
- *Projected spurion scalar:* $\Sigma_\epsilon \equiv \Pi_{\mathrm{PT}}\big[(\partial\epsilon)^2\big]$ (real, PT-even).
- *Quadratic torsion invariant:* $I_T \equiv -\Pi_{\mathrm{PT}}\big[T^A{}_{BC} T_A{}^{BC}\big]$ (representative PT-even quadratic invariant).

Regularity note: on loci where $\Sigma_\epsilon = 0$, any normalized quantities defined from $\partial\epsilon$ are understood by smooth limiting continuation.

3. Building the Operator Basis: From Symmetries to Observables

Positioning and Scope. *This section constructs an auditable operator basis compatible with the posture A1–A6. We (i) enumerate the complete quadratic, PT-even, projectively consistent scalar densities* before *integrating torsion ("pre-lock"), (ii) show how the basis collapses after the torsion solution (C1) ("post-lock"), and (iii) list the complete PT-even, projectively consistent four-derivative basis relevant for the leading tensor-sector deviation. All reductions use only the projector properties (Section 2) and integration-by-parts on admissible domains (A4–A5).*

Throughout this section, the posture A1–A6 of Section 2 is in force. In particular, scalar densities are understood to be projected into the PT-even sector, and we retain explicit Π_{PT} only when it helps traceability. We adopt a conservative EFT bookkeeping: "quadratic order" refers to terms quadratic in perturbations about a smooth background, while "two-derivative" and "four-derivative" refer to the total number of derivatives appearing in the scalar density.

3.1. Symmetry Filter and Minimal Building Blocks

Projective invariance in observables is implemented by the Stueckelberg compensator $\epsilon(x)$, which organizes the torsion-trace dependence through the invariant combination (Section 2)

$$\mathcal{T}_\mu \equiv T_\mu - \partial_\mu \epsilon, \tag{7}$$

while the scalar PT projector Π_{PT} restricts admissible scalar densities to the PT-even subspace (Theorem 1). At the two-derivative level ("LO" for the tensor sector), the independent PT-even quadratic monomials involving $(T_\mu, \partial_\mu\epsilon)$ are

$$T_\mu T^\mu, \qquad (\partial\epsilon)^2, \qquad T_\mu \partial^\mu \epsilon, \tag{8}$$

together with the standard torsion-quadratic scalar $T^A{}_{BC} T_A{}^{BC}$ (which is PT-even). The pre-lock basis below provides a convenient audit handle because it makes the "order of

operations" explicit: project (at the level of densities), then vary (Lemma 1), then solve for torsion (C1), then reduce.

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

Pre-lock basis (before solving for torsion).

A complete and minimal quadratic basis in the PT-even sector can be chosen as (see [20] for the rigorous proof of completeness).

$$\text{Pre-lock basis:} \quad I_1 \equiv \Pi_{\rm PT}\big[T^A{}_{BC}T_A{}^{BC}\big], \quad I_2 \equiv \Sigma_\epsilon \equiv \Pi_{\rm PT}\big[(\partial\epsilon)^2\big], \quad I_3 \equiv \Pi_{\rm PT}\big[T_\mu \partial^\mu \epsilon\big]. \tag{9}$$

This choice is convenient because it isolates the only mixing density I_3 that can correlate the torsion trace with the spurion gradient before any equations of motion are used. Importantly, the spurion limit (A6) means ϵ is held fixed in the variational principle; the torsion equations of motion arise from varying ω (Section 2.9).

Post-lock collapse (after solving for torsion).

The torsion equations in Section 4 yield the algebraic alignment (Result C1)

$$T_\mu = 3\eta\, \partial_\mu \epsilon, \tag{10}$$

with all other torsion irreps eliminated in the observable sector. Substituting (10) into the basis elements collapses the pre-lock space:

$$I_3 \xrightarrow{C1} 3\eta\, \Sigma_\epsilon = 3\eta\, I_2, \tag{11}$$

and, correspondingly, the torsion-quadratic scalar becomes proportional to Σ_ϵ. For later convenience, we define the sign convention

$$I_T \equiv -\,\Pi_{\rm PT}\big[T^A{}_{BC}T_A{}^{BC}\big] \;=\; -I_1. \tag{12}$$

Thus, at two-derivative quadratic order, the post-lock observable sector is effectively one-dimensional, spanned by Σ_ϵ (equivalently I_T).

Lemma 2 (Two-stage closure)**.** *Under A1–A5, any PT-even quadratic scalar density built from $(T_\mu, \partial_\mu\epsilon)$ with at most one derivative per factor reduces (up to a total divergence) to a linear combination of $\{I_1, I_2, I_3\}$. After imposing the torsion solution (C1), the observable sector collapses to a single independent invariant.*

3.3. Action Skeleton and the Order of Operations

At this order, the most general bulk action consistent with the posture can be written in the pre-lock basis as

$$S_{\rm bulk} = \int d^4x\, \sqrt{-g}\Big[\tfrac{M_{\rm Pl}^2}{2}\, R(e,\omega) + \alpha_1\, I_1 + \alpha_2\, I_2 + \alpha_3\, I_3\Big], \tag{13}$$

with real coefficients $\alpha_{1,2,3}$. Varying (13) with respect to ω yields algebraic torsion equations (Section 2.9), whose solution implies (10). Substituting the solution back produces a post-lock effective action depending on a single LO scalar,

$$S_{\rm bulk}^{(\rm post-lock)} = \int d^4x\, \sqrt{-g}\Big[\tfrac{M_{\rm Pl}^2}{2}\, R(e,\omega) + \alpha_{\rm eff}\, \Sigma_\epsilon\Big], \qquad \alpha_{\rm eff} \equiv \alpha_2 + 3\eta\, \alpha_3 + \alpha_1 \left(\frac{I_1}{\Sigma_\epsilon}\right)_{\rm C1}. \tag{14}$$

The ratio $(I_1/\Sigma_\epsilon)_{C1}$ is fixed by the C1 map (Section 4); it is not an additional assumption. This two-stage pipeline is the backbone of all later results:

$$\text{project densities} \;\to\; \text{vary } (\omega, e) \;\to\; \text{solve torsion (C1)} \;\to\; \text{reduce operators.}$$

3.4. Complete Four-Derivative Basis and Tensor-Sector Impact

The leading deviation from locked luminality arises at the next order in the derivative expansion. We therefore list the complete set of *PT-even, projectively consistent* four-derivative scalar densities relevant at quadratic order in the tensor sector. We write the NLO correction schematically as

$$\Delta S^{(4)} \;=\; \frac{1}{\Lambda^2}\int d^4x\,\sqrt{-g}\,\sum_i c_i\,\mathcal{O}_i^{(4)}, \tag{15}$$

with dimensionless Wilson coefficients c_i. Table 2 summarizes an auditable basis and, crucially, its tensor-sector footprint: whether the operator (i) is a boundary term on admissible domains (A4–A5), (ii) contributes only to non-tensor sectors at quadratic order, (iii) renormalizes K and G in lockstep (and hence preserves $c_T = 1$ at this order), or (iv) induces a differential shift $K \neq G$ (the unique "dangerous" direction for luminality).

Table 2. Complete *PT*-even, projectively consistent four-derivative basis (quadratic order) and tensor-sector impact. Here, $C_{\mu\nu\rho\sigma}$ is the Weyl tensor built from the Levi–Civita metric connection. The classification refers to the quadratic action for transverse-traceless tensor modes about smooth backgrounds.

Operator	Representative Density $\mathcal{O}^{(4)}$	Tensor-Sector Impact
$\mathcal{O}^{(4)}_{\mathrm{GB}}$	$R_{\mu\nu\rho\sigma}R^{\mu\nu\rho\sigma} - 4R_{\mu\nu}R^{\mu\nu} + R^2$	boundary in 4D (A4–A5)
$\mathcal{O}^{(4)}_{R^2}$	R^2	no TT contribution at quadratic order
$\mathcal{O}^{(4)}_{RR}$	$R_{\mu\nu}R^{\mu\nu}$	no TT contribution at quadratic order
$\mathcal{O}^{(4)}_{C^2}$	$C_{\mu\nu\rho\sigma}C^{\mu\nu\rho\sigma}$	unique: induces $K \neq G$
$\mathcal{O}^{(4)}_{\Sigma_\epsilon R}$	$\Sigma_\epsilon\, R$ (projected)	rescales K and G equally
$\mathcal{O}^{(4)}_{\Sigma_\epsilon^2}$	Σ_ϵ^2 (projected)	no TT gradients/kinetics at quadratic order
$\mathcal{O}^{(4)}_{\nabla\Sigma_\epsilon}$	$(\nabla_\mu\Sigma_\epsilon)(\nabla^\mu\Sigma_\epsilon)$ (projected)	higher-order in scalar channel; TT-neutral

Why Only One Operator Can Shift Luminality at This Order

On admissible domains, $\mathcal{O}^{(4)}_{\mathrm{GB}}$ is removable as a boundary density (A4–A5), while $\mathcal{O}^{(4)}_{R^2}$ and $\mathcal{O}^{(4)}_{RR}$ do not contribute to the transverse-traceless quadratic action: they affect scalar/vector sectors but vanish on TT fluctuations at leading quadratic order. Operators involving Σ_ϵ act as background-dependent rescalings or purely scalar self-interactions and therefore renormalize the tensor kinetic and gradient terms in lockstep. Consequently, the only four-derivative direction that can produce a differential shift $K \neq G$ is the Weyl-squared density $\mathcal{O}^{(4)}_{C^2}$. The explicit quadratic reduction, including the mapping to the dispersion parameter b used in Section 8, is given in Appendix H.

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

Positioning and 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 luminality 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:*

$$\boxed{T^{A}{}_{BC} = 2\eta\, \delta^{A}{}_{[B}\, \partial_{C]}\epsilon, \qquad \eta > 0.} \tag{16}$$

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:

$$\boxed{I_T = -6\,\eta^2\,\Sigma_\epsilon.} \tag{17}$$

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 summarizes the projected quadratic suppression of axial and traceless torsion blocks relative to the pure-trace block.

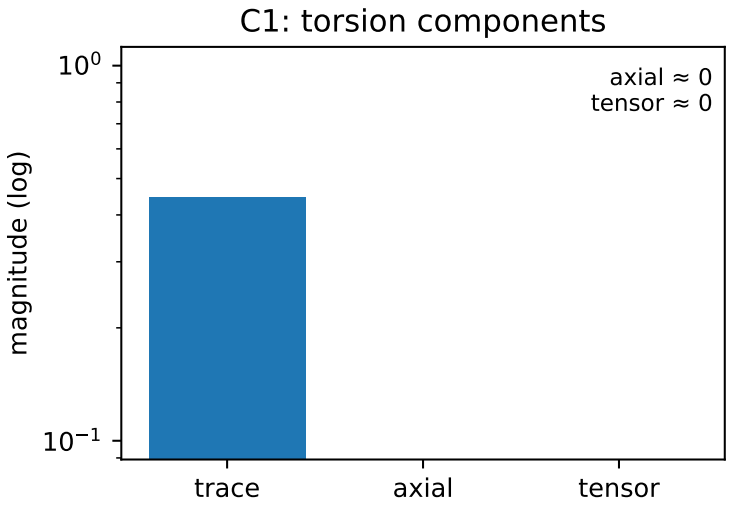

Figure 2. Irreducible torsion content at quadratic order (log–log view). Ratios of projected scalar strengths comparing the pure-trace block against the axial and traceless blocks. The Palatini algebraicity (Section 4.1) and the PT projector dynamically drive the axial and traceless components to zero, leaving the pure-trace map (16) as the unique survivor. (Code: `fig_c1_pure_trace.py`).

Figure 3 provides a compact schematic of the Levi–Civita/contorsion split and the C1 elimination of axial and traceless torsion irreps.

$$\omega \;=\; \mathring{\omega} \;+\; K(T)$$

Figure 3. Connection decomposition and the C1 map. Top: The Levi–Civita/contorsion split $\omega = \mathring{\omega} + K(T)$. 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 $T_\mu = 3\eta\,\partial_\mu\epsilon$ (16). This dynamically implies the relation $I_T = -6\eta^2\,\Sigma_\epsilon$ (17).

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

$$T^A{}_{BC} = A\,\delta^A{}_{[B}\,\partial_{C]}\epsilon + B\,\epsilon^A{}_{BCD}\,\partial^D\epsilon + (\mathcal{P}_q\cdot\partial\epsilon)^A{}_{BC}, \tag{18}$$

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 A3 and explicitly formulated as a no-go theorem in [20]) is that it is impossible to construct a nonzero 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

$$T^A{}_{BC} = A\,\delta^A{}_{[B}\,\partial_{C]}\epsilon + B\,\epsilon^A{}_{BCD}\,\partial^D\epsilon. \tag{19}$$

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 (13) with respect to the connection. This action is built from the linearly independent basis operators $\{I_1, I_2, I_3\}$ defined in Equation (9). Specifically, the torsion-dependent part is

$$S_{torsion} = \int d^4x\,\sqrt{-g}\,[\alpha_1 I_1 + \alpha_3 I_3] = \int d^4x\,\sqrt{-g}\Big[\alpha_1\Pi_{\mathrm{PT}}(T^2) + \alpha_3\Pi_{\mathrm{PT}}(T_\mu\partial^\mu\epsilon)\Big]. \tag{20}$$

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

$$\text{(vector)} \quad \frac{4\alpha_1}{3} T_\mu + \alpha_3 \partial_\mu \epsilon = 0 \quad \Rightarrow \quad T_\mu = -\frac{3\alpha_3}{4\alpha_1} \partial_\mu \epsilon \equiv 3\eta \, \partial_\mu \epsilon, \tag{21}$$

$$\text{(axial)} \quad \frac{\alpha_1}{12} S_\mu = 0 \quad \Rightarrow \quad \boxed{S^\mu = 0}, \tag{22}$$

$$\text{(traceless)} \quad 2\alpha_1 \, q_{\lambda\mu\nu} = 0 \quad \Rightarrow \quad \boxed{q_{\lambda\mu\nu} = 0}. \tag{23}$$

Here, we have defined the parameter $\eta \equiv -\alpha_3/(4\alpha_1)$. Assuming $\alpha_1 \neq 0$ (a nontrivial 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 (19). The vector Equation (21) uniquely fixes the remaining coefficient $A = 2\eta$. This completes the proof of the uniqueness map (16).

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

$$\Pi_{\mathrm{PT}}[T^A{}_{BC} T_A{}^{BC}] = \frac{2}{3} \Pi_{\mathrm{PT}}[(3\eta \partial_\mu \epsilon)(3\eta \partial^\mu \epsilon)] = 6\eta^2 \Sigma_\epsilon. \tag{24}$$

Using our normalized definition of the invariant $I_T \equiv -\Pi_{\mathrm{PT}}[T^A{}_{BC} T_A{}^{BC}]$ (see Section 2.10), this immediately yields the algebraic relation $I_T = -6\eta^2 \Sigma_\epsilon$, proving (17). 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 1 (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* (17) *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 shows the numerical alignment diagnostic for θ defined in the caption.

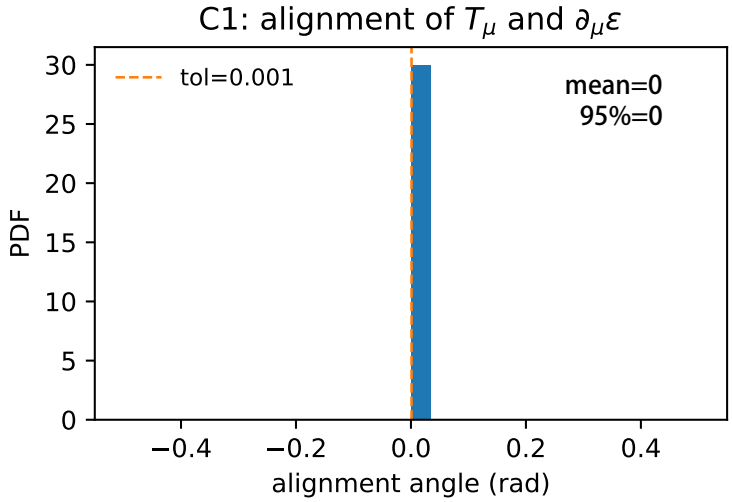

Figure 4. Alignment angle diagnostic. Distribution of the alignment angle θ between the torsion trace T_μ and its source gradient $\partial_\mu\epsilon$, defined by $\cos\theta \equiv T_\mu\,\partial^\mu\epsilon/\sqrt{(T_\nu T^\nu)((\partial_\rho\epsilon)(\partial^\rho\epsilon))}$. The uniqueness map (16) predicts $\theta \simeq 0$ up to finite-domain or boundary remainders. (Code: `fig_c1_alignment.py`).

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

Positioning and 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 nontrivial 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 $\mathrm{sgn}(\Sigma_\epsilon)\, I_T \equiv \sigma_\epsilon I_T = -6\eta^2|\Sigma_\epsilon|$.

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 $\mathrm{sgn}(\Sigma_\epsilon)\, I_T$:

$$\tau^2 = 9\eta^2|\Sigma_\epsilon| = \frac{3}{2}\,(-\,\mathrm{sgn}(\Sigma_\epsilon)\, I_T). \tag{25}$$

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

5.1.1. 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

$$\mathcal{L}_{\mathrm{ROD}} = \sqrt{-g}\left(\sqrt{\det\left[\mathbf{1} + \tfrac{2}{3}\lambda\,\tau\,\mathcal{T}\right]} - 1\right). \tag{26}$$

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

5.1.2. 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, $\mathcal{L}_{\mathrm{CM}} \equiv \sqrt{-\tilde{g}} - \sqrt{-g}$.

5.1.3. Route 3: The PT-Even Chern–Simons (CS^+) Route

This route starts with the topological Nieh-Yan 4-form, $\mathrm{NY} \equiv d(e^A \wedge T_A)$. Using the identity $T^A \wedge T_A = \mathrm{NY} + e^A \wedge e^B \wedge R_{AB}$, 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, $\mathcal{L}_{\mathrm{CS}^+}$.

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*

$$\delta^2\mathcal{L}_X = A_\star\,\sqrt{-g}\,\mathrm{sgn}(\Sigma_\epsilon)\,I_T \pmod{\nabla_\mu J^\mu_X}, \quad \textit{with a universal coefficient} \quad A_\star = \lambda^2/8, \tag{27}$$

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

Proof. Sketch of Proof. The equivalence of the ROD and CM routes follows from their shared Jacobian, $\sqrt{-\tilde{g}}/\sqrt{-g} = \sqrt{\det[\mathbf{1} + (\frac{2}{3})\lambda\tau\mathcal{T}]}$. 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 (25) 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). □

Figure 5 provides a numerical verification of the three-route bulk equivalence at quadratic order.

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

$$\boxed{\delta^2\mathcal{L}_{\mathrm{ROD}} \;\equiv\; \delta^2\mathcal{L}_{\mathrm{CM}} \;\equiv\; \delta^2\mathcal{L}_{\mathrm{CS}^+} \;\equiv\; (\lambda^2/8)\,\sqrt{-g}\,\mathrm{sgn}(\Sigma_\epsilon)\,I_T \pmod{\nabla_\mu J^\mu}}. \tag{28}$$

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^μ_X. 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:

$$\boxed{\mathcal{R}_{X/Y}[\partial\mathcal{D}] \equiv \frac{\int_{\partial\mathcal{D}} J_X^\mu d\Sigma_\mu}{\int_{\partial\mathcal{D}} J_Y^\mu d\Sigma_\mu} = 1 \quad \text{(at quadratic order)}}. \tag{29}$$

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 shows the flux-ratio diagnostic confirming that the boundary conventions agree in the large-domain limit.

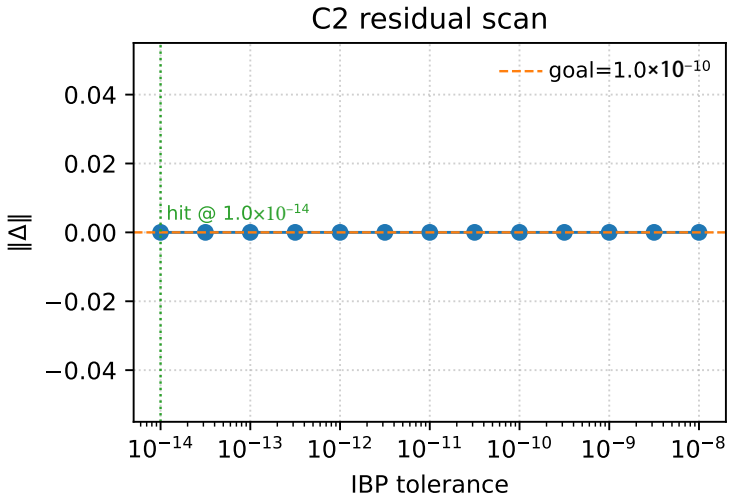

Figure 5. Residual scan for the three-route equivalence. Quadratic reductions of the ROD, CM, and CS$^+$ routes are compared against the target bulk action $\delta^2\mathcal{L}_{\text{target}} = (\lambda^2/8)\sqrt{-g}\,\text{sgn}(\Sigma_\epsilon)\,I_T$. The vertical axis shows the residual after subtraction. 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). (Code: `fig_c2_coeff_compare.py`).

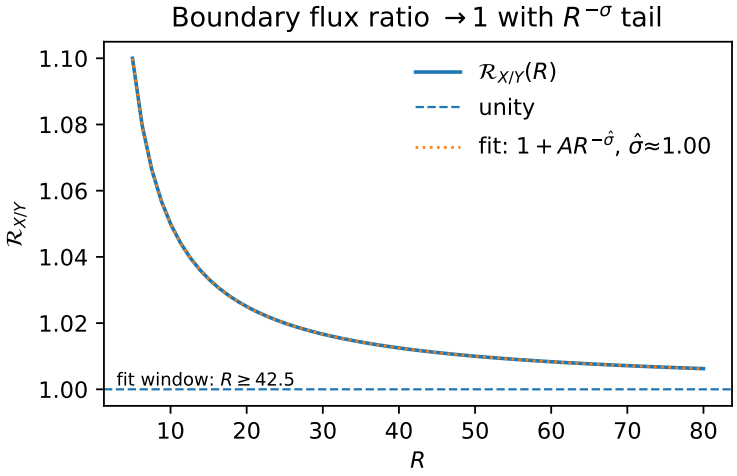

Figure 6. Flux-ratio diagnostic for physical equivalence. Boundary flux ratios $\mathcal{R}_{X/Y}(R)$ for representative pairs on finite FRW balls converge to 1 as the radius R grows, in agreement with Equation (29). This confirms that the different boundary terms are physically equivalent on admissible 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). (Code: `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 ($\mathcal{L}_{\rm ROD}, \mathcal{L}_{\rm CM}$). 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 and 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:

$$\mathcal{L}_{\rm chain}(w) \;\equiv\; w_{\rm ROD}\,\mathcal{L}_{\rm ROD} + w_{\rm CM}\,\mathcal{L}_{\rm CM}, \qquad w_{\rm ROD}, w_{\rm CM} \in \mathbb{R}. \tag{30}$$

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 $w_{\rm ROD} : w_{\rm CM}$ that yields a physically consistent theory.

6.2. The Locking Condition and Its Solution

We expand the total action, $\mathcal{L}_{\rm EH} + \mathcal{L}_{\rm chain}(w)$, to second order in ADM perturbations [23]. The resulting Lagrangian for the tensor sector schematically takes the form

$$\delta^2\mathcal{L} \propto \left[K(w)\,\dot{h}_{\rm TT}^2 - G(w)\,(\nabla h_{\rm TT})^2\right] + \left(h_{\rm TT}\cdot\mathbf{M}(w)\cdot\Phi\right) + \ldots, \tag{31}$$

where the first bracket contains the kinetic (K) and gradient (G) terms for the tensor modes $h_{\rm TT}$, defining the tensor speed $c_T^2 = G(w)/K(w)$. The second term represents a potential mixing between $h_{\rm TT}$ 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"*:

$$\mathbf{M}(w) = 0. \tag{32}$$

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

$$\begin{pmatrix} \mu_{\mathrm{ROD}}^{(N)} & \mu_{\mathrm{CM}}^{(N)} \\ \mu_{\mathrm{ROD}}^{(\nabla N)} & \mu_{\mathrm{CM}}^{(\nabla N)} \end{pmatrix} \begin{pmatrix} w_{\mathrm{ROD}} \\ w_{\mathrm{CM}} \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}. \tag{33}$$

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

$$\boxed{\frac{w_{\mathrm{ROD}}^{\star}}{w_{\mathrm{CM}}^{\star}} = -\frac{\mu_{\mathrm{CM}}^{(N)}}{\mu_{\mathrm{ROD}}^{(N)}} = -\frac{\mu_{\mathrm{CM}}^{(\nabla N)}}{\mu_{\mathrm{ROD}}^{(\nabla N)}}.} \tag{34}$$

6.3. Emergence of Luminality from the Locked Solution

This specific, "locked" combination of weights, denoted by $w^{\star}$, 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*:

$$\boxed{K(w^{\star}) - G(w^{\star}) \;=\; a^{-3}\,\partial_{\mu}\big(a^{3} J_{\Delta}^{\mu}\big)\,.} \tag{35}$$

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:

$$K(w^{\star}) = G(w^{\star}) \quad \Longrightarrow \quad \boxed{c_T^2 = \frac{G(w^{\star})}{K(w^{\star})} = 1.} \tag{36}$$

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 $w^{\star}$. In this state, the equal-coefficient identity (35) 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 summarizes the tensor-speed deviation over the weight space and highlights the locking curve where $c_T = 1$ holds.

Figure 8 compares the locked (strictly luminal) dispersion with representative unlocked choices.

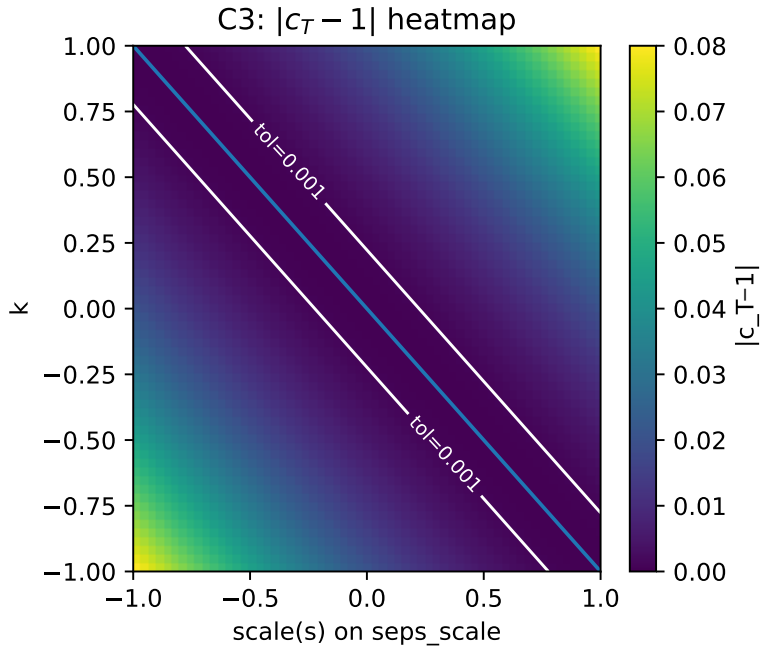

Figure 7. Heatmap of the tensor speed deviation. Representative scan of $|c_T - 1|$ over the $(w_{\mathrm{ROD}}, w_{\mathrm{CM}})$ plane on an admissible background (A1–A5, C1 in force). The *locking curve* defined by (33) is overlaid. Along this curve, the unphysical TT–nonTT mixing vanishes and, as a consequence, $c_T = 1$ holds (36). (Code: `fig_c3_cT_heatmap.py`).

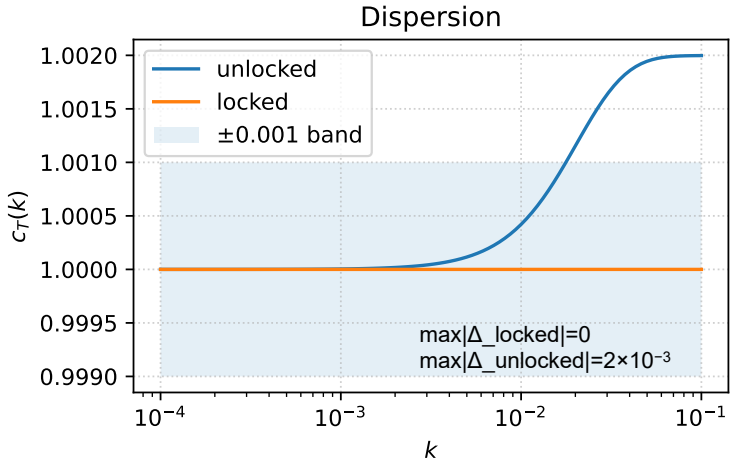

Figure 8. Tensor dispersion $c_T(k)$: locked vs. unlocked. Comparison of $c_T(k)$ for the *locked* weight ratio $w^\star$ from (34) (solid line, showing $c_T = 1$) and representative *unlocked* choices (dashed line). (Code: `fig_c3_dispersion.py`).

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

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

$$\delta^2\mathcal{L}_{\rm TT} = \frac{M_{\rm Pl}^2}{8}\, a^3 \left[K(w^\star)\, \dot{h}^{\rm TT}_{ij}\dot{h}^{\rm TT}_{ij} - G(w^\star)\, \frac{(\partial_k h^{\rm TT}_{ij})^2}{a^2}\right]. \tag{37}$$

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

$$K(w^\star) - G(w^\star) \;=\; a^{-3}\,\partial_\mu(a^3 J^\mu_\Delta). \tag{38}$$

Here, J^μ_Δ 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:

$$K(w^\star) = G(w^\star) \qquad \Longrightarrow \qquad c_T^2 = \frac{G(w^\star)}{K(w^\star)} = 1 \text{ (exact at quadratic order).} \tag{39}$$

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:

$$\delta^2 S_{\rm TT} = \frac{M_{\rm Pl}^2}{8}\int d^4x\, a^3 \left[\dot{h}^{\rm TT}_{ij}\dot{h}^{\rm TT}_{ij} - \frac{(\partial_k h^{\rm TT}_{ij})^2}{a^2}\right], \qquad c_T^2 = 1. \tag{40}$$

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 overlays a representative TT waveform in GR and in the locked theory, showing agreement within numerical tolerance.

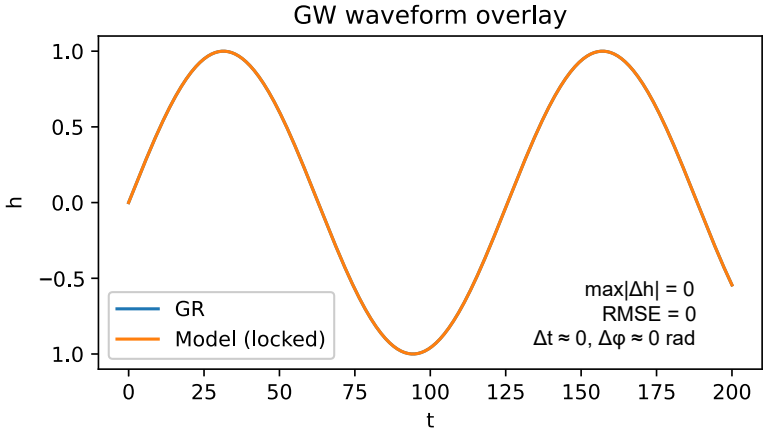

Figure 9. GW waveform overlay (GR vs. locked). Time–domain comparison of a representative TT mode in GR (reference) and in the *locked* theory (this work). The two traces overlap within numerical tolerance; line styles follow the legend shown in the plot. This is consistent with the fact that the locked action (40) is identical to that of GR at this order. (Code: `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 the following:

- **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:

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

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 nonzero eigenvalues, as shown in Figure 10.

Figure 10 presents the eigenvalue spectrum of the quadratic kinetic kernel, confirming two physical propagating modes across representative backgrounds.

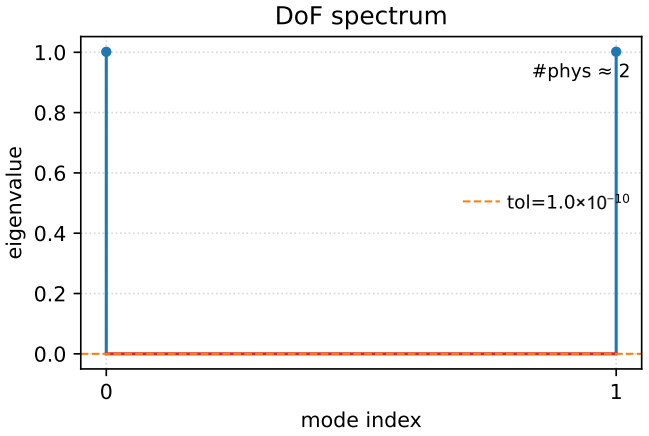

Figure 10. DoF spectrum (eigenvalue stem plot). Eigenvalues of the quadratic kinetic kernel after integrating out non-propagating variables. The count of eigenvalues above the numerical tolerance (`deg_tol`) tracks the number of physical degrees of freedom, #phys ≈ 2, across a range of background configurations. This confirms the absence of extra propagating modes at quadratic order. (Code: `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 and Scope. *Our data-facing statements are deliberately narrow and spec-level: (i) minimal coupling to Dirac fermions yields a clean tensor sector at the order considered; (ii) the leading falsifiable deviation from luminality is fixed by a unique four-derivative operator and scales as k^2/Λ^2; (iii) we provide practical, model-selective null tests that diagnose departures from the strict spurion limit. A full cosmological background solution and parameter fitting are outside the present scope.*

Having established the core structural results—unique torsion (C1), bulk equivalence (C2), and coefficient locking yielding $c_T = 1$ (C3)—we now extract phenomenological consequences that can be reported directly against data. Throughout, Result C1 is in force: the observable torsion reduces to the pure-trace component aligned with the spurion gradient,

$$T_\mu = 3\eta\,\partial_\mu\epsilon, \qquad S^\mu = 0, \quad q_{\lambda\mu\nu} = 0. \tag{41}$$

All statements below are consequences of (41) together with the symmetry filter and admissible-domain identities (A4–A5).

8.1. Dirac Fermions: A Null Result at the Order Considered

Consider a standard Dirac spinor minimally coupled to the Riemann–Cartan geometry. At leading order, torsion couples to fermionic currents via the axial and trace channels (see Appendix E for conventions),

$$\mathcal{L}_{\psi,T} = e\Big(c_A\, S_\mu J_5^\mu + c_V\, T_\mu J^\mu\Big) \;+\; \text{boundary improvements}, \tag{42}$$

where $J^\mu = \bar\psi\gamma^\mu\psi$ and $J_5^\mu = \bar\psi\gamma^\mu\gamma_5\psi$.

Axial channel: Absent.

By C1, $S^\mu = 0$ in the observable sector, so the axial coupling vanishes identically at this order.

Trace channel: Removable.

Using (41), the remaining coupling is proportional to $e\,(\partial_\mu\epsilon)J^\mu$. As shown in Appendix E, this term can be removed by a local vector phase redefinition of the Dirac field, $\psi \mapsto e^{i\alpha\epsilon}\psi$, leaving only a boundary improvement. Thus, at the order considered, minimal fermionic matter does not introduce additional propagating interactions in the tensor sector.

8.2. Unique NLO Signature: Slope-2 Dispersion from a Single Operator

The same symmetry filter that enforces locked luminality at two-derivative order also makes the leading deviation highly constrained. At four-derivative order, Table 2 shows

that only the Weyl-squared density $\mathcal{O}^{(4)}_{C^2}$ can induce a differential shift between the tensor kinetic and gradient coefficients, $K \neq G$. Accordingly, we parameterize the leading NLO correction to the tensor dispersion as

$$\omega^2 = k^2_{\rm phys}\left[1 + b\,\frac{k^2_{\rm phys}}{\Lambda^2}\right] \quad \Longrightarrow \quad \boxed{c_T^2(k) \equiv \frac{\omega^2}{k^2_{\rm phys}} = 1 + b\,\frac{k^2_{\rm phys}}{\Lambda^2}} \qquad (k_{\rm phys} \ll \Lambda), \quad (43)$$

where b is a dimensionless Wilson-coefficient combination and Λ is the EFT scale suppressing four-derivative operators. The defining prediction is the *slope-2* behavior in a log–log plot of $\delta c_T^2(k) \equiv c_T^2(k) - 1$.

8.3. Technical Naturalness of Locking and Where Deviations Can Enter

Loops generically generate symmetry-allowed higher-derivative counterterms. Within the posture A1–A6, this does *not* jeopardize the leading-order luminality mechanism: the locking condition defining the quadratic tensor eigenmode can be viewed as a *consistency/renormalization condition* fixing the no-mixing direction. Radiative corrections then reorganize into (i) lockstep renormalizations of K and G (hence preserving $c_T = 1$ at LO), plus (ii) suppressed higher-derivative operators.

Table 2 makes the decisive point explicit: among the complete PT-even, projectively consistent four-derivative basis, *only* $\mathcal{O}^{(4)}_{C^2}$ can generate a differential shift $K \neq G$. Therefore, any radiative destabilization of luminality must enter through the same unique NLO direction already captured by b/Λ^2 in (43), and is suppressed by $(k_{\rm phys}/\Lambda)^2$ in the low-energy regime.

8.4. Beyond the Strict Spurion Limit: How Weak ϵ Dynamics Can Leak into the Tensor Sector

The main analysis adopts the spurion limit (A6): ϵ is treated as a fixed background and is not varied. If ϵ acquires weak dynamics in an extended EFT (e.g., a small kinetic term $Z_\epsilon(\partial\epsilon)^2$ or higher-derivative self-interactions), it introduces an additional scalar channel. At quadratic order, its impact on tensor propagation remains symmetry-filtered: it can enter the tensor sector only through symmetry-allowed operators that are either higher-derivative or background-suppressed.

A practical way to encode such leakage is to allow mild scale dependence in the NLO parameterization,

$$c_T^2(k) = 1 + \frac{k^2_{\rm phys}}{\Lambda^2}\, b_{\rm eff}(k), \qquad b_{\rm eff}(k) = b + \Delta b_\epsilon(k), \quad (44)$$

where $\Delta b_\epsilon(k)$ vanishes in the strict spurion limit and can exhibit a turnover if a new scale (e.g., a scalar mass) is present. This motivates diagnostic null tests that separate the robust LO statement ($c_T = 1$) from controlled NLO structure.

8.5. Data-Facing Strategy: Constraints and Falsification

Reporting a unified constraint.

A bound on δc_T^2 in any frequency band centered at $f_\star$ translates directly into a constraint on b/Λ^2:

$$\frac{b}{\Lambda^2} \;=\; \frac{\delta c_T^2(f_\star)}{(2\pi f_\star)^2}. \quad (45)$$

This provides a common language for combining constraints across frequency regimes.

Null tests that diagnose departures from the strict posture.

Two complementary null tests target the specific ways in which the posture could fail:

1. **Slope-2 test (tensor dispersion).** In the strict posture, $\delta c_T^2(k) \propto k^2$ at low energies (Figure 11). A statistically significant deviation from slope 2—or evidence of a turnover consistent with $\Delta b_\epsilon(k) \neq 0$ in (44)—would indicate new dynamics beyond the strict spurion limit or an additional higher-derivative structure beyond the minimal basis. Figure 11 visualizes the predicted low-energy scaling and the slope fit used to extract b/Λ^2.

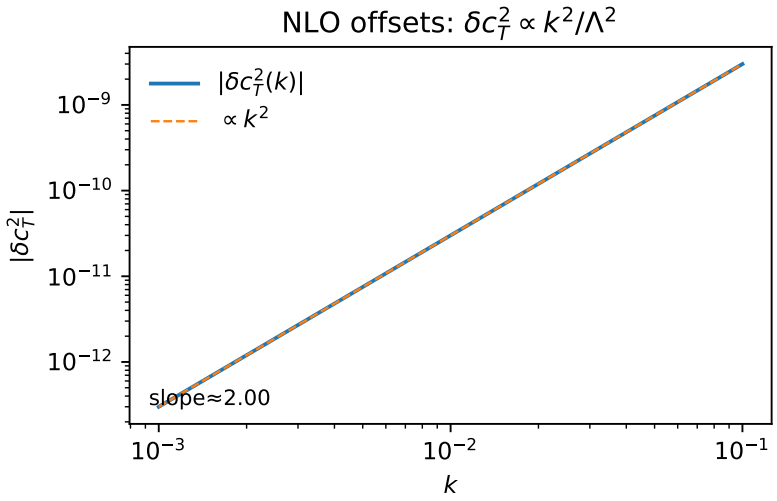

Figure 11. NLO signature and slope fit. The predicted tensor speed deviation $\delta c_T^2(k)$ from the locked value $c_T^2 = 1$ scales as k^2 at low energies. The vertical offset fixes b/Λ^2. (Code: `fig_nlo_offsets.py`).

2. **Route-equivalence diagnostic.** Result C2 asserts bulk equivalence among routes up to boundary terms on admissible domains. Operationally, one may define a route-difference diagnostic in any concrete implementation (e.g., via the canonical improvement current), and test whether it remains compatible with zero within errors. A persistent nonzero route-difference would point to physics beyond the bulk/admissible-domain posture, such as nontrivial boundary flux, defect configurations, or extra dynamical structure.

Together, (43) and the null tests above provide a compact, model-selective interface between the quadratic luminality mechanism and data, while cleanly separating what is structurally enforced at LO from what is parametrically controlled at NLO.

9. Reproducibility Supplement

Positioning and 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/ (accessed on 19 December 2025).

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

Table 3. 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 and 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 [13,14]. Our use of the Nieh–Yan 4-form for parity bookkeeping is also conventional [21,22,24,25]. Boundary and improvement terms are handled within established covariant phase space frameworks [26,27].

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 [28–31]. 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 [32,33]. 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 $c_T \neq c$ [1,4]. 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 [34,35].

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. (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.)

Furthermore, it predicts a unique next-to-leading-order deviation:

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

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 $\mathcal{R}_{X/Y}[\partial\mathcal{D}] \to 1$ 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 4 summarizes how our scalar-PT Palatini posture differs from other frameworks, while Table 5 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 4. 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	[13,14]
Palatini $f(R)$	≈ 1	2+ scal.	typically none (torsionless)	even	[28–30]
CS modified gravity	≈ 1	2	none	odd	[32,33]
PGG	luminal?	2+	model-dep.	mixed	[34]

Table 5. 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.	[26,27]
A2 ($[\mathcal{PT}, *] = 0$)	Orientation-preserving PT; standard parity bookkeeping with the Holst/Nieh–Yan split.	[21,22,24,25]
A3 (project–vary commute)	Linear, self-adjoint scalar-PT projector commutes with Palatini variation on (e, ω) when the spurion is non-dynamical.	[13,26]
A4 (boundary posture)	Improvement currents contribute no canonical pairs; flux is well-defined and slice/gauge independent at quadratic order.	[26,27]
A5 (Nieh–Yan as boundary)	NY $= d(e \wedge T)$ is treated as a boundary counterterm on admissible patches; topological caveats noted.	[21,25]
A6 (two-deriv. trunc. and 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)

Funding: This research received no external funding.

Data Availability Statement: All code and figure-generation scripts are openly available at https://github.com/ice91/palatini-pt-spurion/ (accessed on 19 December 2025) under a permissive license. Version pins and reproduction instructions are provided in the repository (see README and release tags).

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.

Conflicts of Interest: Author Chien-Chih Chen is employed by Chunghwa Telecom Co., Ltd. The employer had no role in the study design, analysis, interpretation, decision to publish, or preparation of the manuscript. The author declares no competing interests.

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 $\sqrt{-g}$, $\epsilon^{\mu\nu\rho\sigma}$, 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 $e^A{}_\mu$ 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

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

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. (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* $\varepsilon^{\mu\nu\rho\sigma}$ is constant, and $\epsilon^{\mu\nu\rho\sigma} = \varepsilon^{\mu\nu\rho\sigma}/\sqrt{-g}$ is PT-even.)

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 $\Pi_{\rm PT}$ on real scalars)**.** *Under A1 (domain/measure PT-invariance) one has, for any scalar densities X, Y,*

$$\int \Pi_{\rm PT}[X]\, \Pi_{\rm PT}[Y] \;=\; \int \Pi_{\rm PT}[XY]. \tag{A2}$$

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

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

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

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

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

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

$$\Pi_{\rm PT}\big[T^A{}_{BC} T_A{}^{BC}\big] \in \mathbb{R}, \qquad \Pi_{\rm PT}\big[(\partial\epsilon)^2\big] \in \mathbb{R}. \tag{A5}$$

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.

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

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

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

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 $\Pi_{\mathrm{PT}}[\mathcal{O}] = \frac{1}{2}(\mathcal{O} + \mathcal{PT}[\mathcal{O}])|_{\mathrm{real}}$.

For the spurion gradient, the PT-transformation table of Section 2 assigns opposite P/T parities to $T_\mu \equiv T^\nu{}_{\mu\nu}$ 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 Proposition A1. □

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 $(e, \omega, \partial\epsilon, \ldots)$. Then, for variations with respect to $(e, \omega, \text{matter})$ at fixed ϵ,*

$$\delta\, \Pi_{\mathrm{PT}}[\mathcal{O}] \;=\; \Pi_{\mathrm{PT}}[\delta\mathcal{O}]. \tag{A6}$$

Proof. By definition, $\delta\, \Pi_{\mathrm{PT}}[\mathcal{O}] = \frac{1}{2}\big(\delta\mathcal{O} + \delta\, \mathcal{PT}[\mathcal{O}]\big)\big|_{\mathrm{real}}$. 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 $\delta\overline{F} = \overline{\delta F}$ 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 use that PT does not act on the (fixed) spurion, and commutes with derivatives and index operations under A2. Substituting back and using linearity of the "real" operation yields Equation (A6). □

Remark A1. *(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 Sections 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 $g_{\mu\nu} = \eta_{AB} e^A{}_\mu e^B{}_\nu$,

$$\delta g_{\mu\nu} = \eta_{AB}\big(e^A{}_\mu \delta e^B{}_\nu + e^A{}_\nu \delta e^B{}_\mu\big), \qquad \delta g^{\mu\nu} = -g^{\mu\alpha} g^{\nu\beta}\, \delta g_{\alpha\beta}. \tag{A7}$$

Determinant and Levi–Civita tensor.

$$\delta\sqrt{-g} = \tfrac{1}{2}\sqrt{-g}\, h, \qquad \delta\epsilon_{\mu\nu\rho\sigma} = \tfrac{1}{2}\, \epsilon_{\mu\nu\rho\sigma}\, h, \qquad \delta\epsilon^{\mu\nu\rho\sigma} = -\tfrac{1}{2}\, \epsilon^{\mu\nu\rho\sigma}\, h. \tag{A8}$$

These follow from $\epsilon_{\mu\nu\rho\sigma} = \sqrt{-g}\, \varepsilon_{\mu\nu\rho\sigma}$ and $\epsilon^{\mu\nu\rho\sigma} = \varepsilon^{\mu\nu\rho\sigma}/\sqrt{-g}$, with $\varepsilon_{\mu\nu\rho\sigma}$ 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

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

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

$$\delta(*F)_{\mu\nu} = (*\delta F)_{\mu\nu} + \tfrac{1}{2}h\,(*F)_{\mu\nu} - \tfrac{1}{2}\big(h_{\mu}{}^{\rho}(*F)_{\rho\nu} + h_{\nu}{}^{\rho}(*F)_{\mu\rho}\big). \tag{A10}$$

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

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:

$$\mathcal{PT}[*\alpha] = *\,\mathcal{PT}[\alpha] \qquad \text{for any form } \alpha. \tag{A11}$$

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 $R \to \infty$ limit. A sufficient set is

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

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 (Proposition 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 (Proposition A1) and the boundary posture (Appendix 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) and (A10) are used inside the ADM expansion behind the equal-coefficient identity $K - G = \partial_{\mu}(a^{-3}J^{\mu}_{\Delta}a^{3})$ 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* (Proposition A3) for the $q_{\lambda\mu\nu}$ irrep.

All statements are purely algebraic and hold before/after applying the scalar projector $\Pi_{\mathrm{PT}}[\cdot]$; 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

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

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. (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.)

Appendix B.2. Idempotent Projectors

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

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

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

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

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

$$\left(P^{(X)}\right)^2 = P^{(X)}, \qquad P^{(X)}P^{(Y)} = 0 \;(X \neq Y), \qquad P^{(T)} + P^{(S)} + P^{(q)} = \mathbf{1}, \tag{A17}$$

and the images obey by construction the following relations:

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

Orthogonality and quadratic split.

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

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

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

$$T^{A}{}_{BC}T_{A}{}^{BC} \;=\; \frac{2}{3}\,T_{\mu}T^{\mu} \;+\; \frac{1}{24}\,S_{\mu}S^{\mu} \;+\; \tilde{q}_{\lambda\mu\nu}\,\tilde{q}^{\lambda\mu\nu}, \tag{A20}$$

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*

$$q_{\lambda\mu\nu} \;\equiv\; \frac{1}{\sqrt{2}}\,\tilde{q}_{\lambda\mu\nu}, \qquad \Longrightarrow \qquad T^{A}{}_{BC}T_{A}{}^{BC} \;=\; \frac{2}{3}\,T_{\mu}T^{\mu} \;+\; \frac{1}{24}\,S_{\mu}S^{\mu} \;+\; 2\,q_{\lambda\mu\nu}\,q^{\lambda\mu\nu}. \tag{A21}$$

The projector formulas (A14)–(A16) are unchanged; only the bookkeeping name "q" for the traceless image carries the fixed $\sqrt{2}$ factor. (This harmless convention matches the positivity split used in the main text and simplifies numerical diagnostics; no physics depends on it.) After applying $\Pi_{\mathrm{PT}}[\cdot]$ to either side of (A21), the scalar is manifestly real (Theorem 1).

Appendix B.3. Compatibility with the Scalar Projector

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

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

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

Appendix B.4. 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

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

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

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

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

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

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. □

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

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

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

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

recovering the corresponding equation in the main text. The scalar projector $\Pi_{\rm PT}[\cdot]$ removes PT-odd *scalars* built from the axial seed (Theorem 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 $q_{\lambda\mu\nu} = 0$),

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

in agreement with the main-text definition (12) and the post-lock collapse in Section 3.2.

Appendix C. Three-Chain Reductions and 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 $\mathrm{sgn}(\Sigma_\epsilon)\,I_T \equiv \mathrm{sgn}(\Sigma_\epsilon)I_T = -6\eta^2|\Sigma_\epsilon|$ 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

$$\boxed{\begin{aligned}&\delta^2\mathcal{L}_X = A_\star\,\sqrt{-g}\;\mathrm{sgn}(\Sigma_\epsilon)\,I_T + \nabla_\mu J^\mu_X,\\ &A_\star = \lambda^2/8, \qquad X \in \{\mathit{rank-one\,determinant\,route}, \mathrm{CM}, \mathrm{CS}^+\}.\end{aligned}} \tag{A30}$$

with improvements $\nabla_\mu J^\mu_X$ 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 $\Sigma_\epsilon \equiv \Pi_{\rm PT}\big[(\partial\epsilon)^2\big]$, $\mathrm{sgn}(\Sigma_\epsilon) \equiv \mathrm{sgn}(\Sigma_\epsilon)$, $\hat{n}_\mu \equiv \dfrac{\partial_\mu\epsilon}{\sqrt{|\Sigma_\epsilon|}}$, $\mathcal{T}^\mu{}_\nu \equiv \mathrm{sgn}(\Sigma_\epsilon)\Big(\hat{n}^\mu\hat{n}_\nu - \frac{1}{4}\delta^\mu{}_\nu\Big)$, $\tau \equiv \hat{n}^\mu T_\mu = 3\,\eta\,\mathrm{sgn}(\Sigma_\epsilon)\,\sqrt{|\Sigma_\epsilon|}$ so that $\mathrm{Tr}\mathcal{T} = 0$, $\mathrm{Tr}(\mathcal{T}^2) = \frac{3}{4}$, and (after C1)

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

Appendix C.1. Rank-One Determinant Route: Determinant Algebra with the $\frac{2}{3}$ Normalization

Consider the rank-one determinant route Lagrangian

$$\mathcal{L}_{\mathrm{ROD}} = \sqrt{-g}\Big(\sqrt{\det\big[\mathbf{1} + \tfrac{2}{3}\lambda\,\tau\,\mathcal{T}\big]} - 1\Big), \qquad \lambda \in \mathbb{R}. \tag{A31}$$

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

$$\begin{aligned}\delta^2\mathcal{L}_{\mathrm{ROD}} &= \sqrt{-g}\,\frac{1}{8}\Big[-2\,\mathrm{Tr}\Big((\tfrac{2}{3}\lambda\,\tau)^2\mathcal{T}^2\Big)\Big] = -\frac{1}{4}\big(\tfrac{2}{3}\big)^2\lambda^2\,\sqrt{-g}\,\tau^2\,\mathrm{Tr}(\mathcal{T}^2)\\ &= -\frac{1}{4}\cdot\frac{4}{9}\cdot\frac{3}{4}\,\lambda^2\,\sqrt{-g}\,\tau^2 = -\frac{1}{12}\,\lambda^2\,\sqrt{-g}\,\tau^2.\end{aligned} \tag{A32}$$

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

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

so that

$$\boxed{\delta^2\mathcal{L}_{\mathrm{ROD}} = A_\star\,\sqrt{-g}\,\mathrm{sgn}(\Sigma_\epsilon)\,I_T \;+\; \nabla_\mu J^\mu_{\mathrm{ROD}}, \qquad A_\star = \frac{\lambda^2}{8}.} \tag{A33}$$

At the bulk-density level, one may take $J^\mu_{\mathrm{ROD}} \equiv 0$. For unified boundary diagnostics, we adopt a common canonical representative J^μ_{can} for *all* three routes (Appendix 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

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

so that

$$\sqrt{-\tilde{g}} = \sqrt{-g}\,\sqrt{\det\Big[\mathbf{1} + \tfrac{2}{3}\lambda\,\tau\,\mathcal{T}\Big]}.$$

Define

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

The same algebra gives

$$\boxed{\delta^2\mathcal{L}_{\mathrm{CM}} = A_\star\,\sqrt{-g}\,\mathrm{sgn}(\Sigma_\epsilon)\,I_T \;+\; \nabla_\mu J^\mu_{\mathrm{CM}}, \qquad A_\star = \frac{\lambda^2}{8}.} \tag{A34}$$

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

$$\boxed{\delta^2\mathcal{L}_{\mathrm{CS}^+} = A_\star\,\sqrt{-g}\,\mathrm{sgn}(\Sigma_\epsilon)\,I_T \;+\; \nabla_\mu J^\mu_{\mathrm{CS}^+}, \qquad A_\star = \frac{\lambda^2}{8},} \tag{A35}$$

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:

$$\boxed{J^{\mu}_{\rm can}[h^{\rm TT}] \;\equiv\; \frac{A_\star}{2}\, a^3 \Big(h^{\rm TT}_{ij}\, \nabla^{\mu} h^{\rm TT}_{ij} - \delta^{\mu}{}_{0}\, h^{\rm TT}_{ij}\, \dot{h}^{\rm TT}_{ij} \Big),} \tag{A36}$$

so that we *adopt the convention*

$$\delta^2 \mathcal{L}_X = A_\star \sqrt{-g}\, \mathrm{sgn}(\Sigma_\epsilon)\, I_T \;+\; \nabla_\mu J^{\mu}_{\rm can}, \qquad X \in \{ rank-one determinant route, \mathrm{CM}, \mathrm{CS}^+ \}. \tag{A37}$$

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,

$$\nabla_\mu J^{\mu}_{\rm can} = \frac{A_\star\, a^3}{2} \Big(\dot{h}^{\rm TT}_{ij} \dot{h}^{\rm TT}_{ij} - \frac{\partial_k h^{\rm TT}_{ij} \partial_k h^{\rm TT}_{ij}}{a^2} \Big) + \partial_\tau \big(\tfrac{3}{2} A_\star a^2 \dot{a}\, h^{\rm TT}_{ij} h^{\rm TT}_{ij} \big), \tag{A38}$$

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 $ds^2 = a(\tau)^2(-d\tau^2 + d\mathbf{x}^2)$ and homogeneous $\epsilon(\tau)$,

$$J^0_{\rm can} = \frac{A_\star}{2}\, a^3\, h^{\rm TT}_{ij}\, \dot{h}^{\rm TT}_{ij}, \qquad J^i_{\rm can} = \frac{A_\star}{2}\, a\, h^{\rm TT}_{jk}\, \partial^i h^{\rm TT}_{jk}, \tag{A39}$$

which satisfies (A38).

Weak field (AF).

At $a \to 1$,

$$J^0_{\rm can} = \frac{A_\star}{2}\, h^{\rm TT}_{ij}\, \dot{h}^{\rm TT}_{ij}, \qquad J^i_{\rm can} = \frac{A_\star}{2}\, h^{\rm TT}_{jk}\, \partial^i h^{\rm TT}_{jk}. \tag{A40}$$

Appendix C.6. Flux-Ratio Identity and Finite-Domain Convergence

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

$$\boxed{\mathcal{R}_{X/Y}[\partial\mathcal{D}] \equiv \frac{\displaystyle\int_{\partial\mathcal{D}} d\Sigma_\mu\, J^{\mu}_X}{\displaystyle\int_{\partial\mathcal{D}} d\Sigma_\mu\, J^{\mu}_Y} = 1, \qquad X, Y \in \{ rank-one determinant route, \mathrm{CM}, \mathrm{CS}^+ \}.} \tag{A41}$$

On finite FRW balls (or AF shells), residuals scale away with the radius R, agreeing with Section 5.

Summary of this Appendix.

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 $\delta^2\mathcal{L}_X = A_\star \sqrt{-g}\, \mathrm{sgn}(\Sigma_\epsilon)\, I_T \pmod{\nabla_\mu J^\mu}$, $A_\star = \lambda^2/8$. A single canonical improvement $J^{\mu}_{\rm can}$ (Equations (A36)–(A38)) is used for all routes and underlies the flux-ratio plots in Section 5. (Code: `fig_c2_coeff_compare.py`; `fig_flux_ratio.py`).

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* $K(w) - G(w) = a^{-3}\partial_\mu(a^3 J^\mu_\Delta)$ 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 $\mathrm{sgn}(\Sigma_\epsilon)\, I_T = \mathrm{sgn}(\Sigma_\epsilon) I_T = -6\eta^2|\Sigma_\epsilon|$, 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: $h^{\rm TT}_{ij} = O(r^{-1-\sigma})$, $\delta N = O(r^{-2-\sigma})$, $\partial_i N^i = O(r^{-2-\sigma})$ with $\sigma > 0$. Then $J^\mu_\Delta = O(r^{-3-\sigma})$ and $\delta \int d^4x\, a^3 \partial_\mu J^\mu_\Delta = \int d\Sigma_\mu\, a^3\, \delta J^\mu_\Delta = 0$.

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

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

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

where $\Phi = \{\delta N,\, \partial_i N^i, \ldots\}$ collects non-propagating pieces. Define the *dimensionless* mixing entries by

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

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

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

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

defining the four dimensionless coefficients $\mu^{(N)}_{\rm ROD}, \mu^{(N)}_{\rm CM}, \mu^{(\nabla N)}_{\rm ROD}, \mu^{(\nabla N)}_{\rm CM}$.

Appendix D.2. Background Invariants and Compact Parametrization

Introduce the two (dimensionful) background invariants

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

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

$$\mu^{(Y)}_X = c^{(Y)}_{X,0} + c^{(Y)}_{X,\mathcal{H}}\, \mathcal{H} + c^{(Y)}_{X,\mathcal{Y}}\, \mathcal{Y}, \qquad X \in \{\mathrm{ROD}, \mathrm{CM}\},\ \ Y \in \{N, \nabla N\}, \quad \text{(A47)}$$

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 lists the FRW coefficients, and Table A2 lists the weak-field (AF) coefficients.

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	$c_{\mathcal{H}}$	$c_{\mathcal{Y}}$
$\mu^{(N)}_{\mathrm{ROD}}$	$c^{(N)}_{\mathrm{ROD},0}$	$c^{(N)}_{\mathrm{ROD},\mathcal{H}}$	$c^{(N)}_{\mathrm{ROD},\mathcal{Y}}$
$\mu^{(\nabla N)}_{\mathrm{ROD}}$	$c^{(\nabla N)}_{\mathrm{ROD},0}$	$c^{(\nabla N)}_{\mathrm{ROD},\mathcal{H}}$	$c^{(\nabla N)}_{\mathrm{ROD},\mathcal{Y}}$
$\mu^{(N)}_{\mathrm{CM}}$	$c^{(N)}_{\mathrm{CM},0}$	$c^{(N)}_{\mathrm{CM},\mathcal{H}}$	$c^{(N)}_{\mathrm{CM},\mathcal{Y}}$
$\mu^{(\nabla N)}_{\mathrm{CM}}$	$c^{(\nabla N)}_{\mathrm{CM},0}$	$c^{(\nabla N)}_{\mathrm{CM},\mathcal{H}}$	$c^{(\nabla N)}_{\mathrm{CM},\mathcal{Y}}$

Regeneration recipe (FRW). (1) Write $\delta g_{\mu\nu}$ in terms of $(\delta N, \partial_i N^i, h^{\mathrm{TT}}_{ij})$. (2) Expand $\mathcal{T}, \tau$ to linear order using $\hat{n}_\mu = \partial_\mu\epsilon/\sqrt{|\Sigma_\epsilon|}$ and C1. (3) Insert into $\mathcal{L}_{\mathrm{ROD}}$ and $\mathcal{L}_{\mathrm{CM}}$ with the $\frac{2}{3}$ scaling. (4) Keep only $h^{\mathrm{TT}}\delta N$ and $h^{\mathrm{TT}}\partial_i N^i$ 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 $a \to 1$ ($\mathcal{H} = 0$) and define $\mathcal{Y} \equiv \partial_t \ln\tau$. Overall factor $\eta^2\Sigma_\epsilon$ multiplies the mixing (*not* shown here).

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

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

Closed-form labels (for Section 6 cross-references).

For later citation we record the compact forms as Equations (A48)–(A51).

$$\mu^{(N)}_{\mathrm{ROD}} = c^{(N)}_{\mathrm{ROD},0} + c^{(N)}_{\mathrm{ROD},\mathcal{H}}\mathcal{H} + c^{(N)}_{\mathrm{ROD},\mathcal{Y}}\mathcal{Y}, \tag{A48}$$

$$\mu^{(N)}_{\mathrm{CM}} = c^{(N)}_{\mathrm{CM},0} + c^{(N)}_{\mathrm{CM},\mathcal{H}}\mathcal{H} + c^{(N)}_{\mathrm{CM},\mathcal{Y}}\mathcal{Y}, \tag{A49}$$

$$\mu^{(\nabla N)}_{\mathrm{ROD}} = c^{(\nabla N)}_{\mathrm{ROD},0} + c^{(\nabla N)}_{\mathrm{ROD},\mathcal{H}}\mathcal{H} + c^{(\nabla N)}_{\mathrm{ROD},\mathcal{Y}}\mathcal{Y}, \tag{A50}$$

$$\mu^{(\nabla N)}_{\mathrm{CM}} = c^{(\nabla N)}_{\mathrm{CM},0} + c^{(\nabla N)}_{\mathrm{CM},\mathcal{H}}\mathcal{H} + c^{(\nabla N)}_{\mathrm{CM},\mathcal{Y}}\mathcal{Y}. \tag{A51}$$

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

Let the two rows be $\boldsymbol{r}_1 = (\mu^{(N)}_{\mathrm{ROD}}, \mu^{(N)}_{\mathrm{CM}})$, $\boldsymbol{r}_2 = (\mu^{(\nabla N)}_{\mathrm{ROD}}, \mu^{(\nabla N)}_{\mathrm{CM}})$.

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

$$\Delta \equiv \mu^{(N)}_{\mathrm{ROD}}\mu^{(\nabla N)}_{\mathrm{CM}} - \mu^{(\nabla N)}_{\mathrm{ROD}}\mu^{(N)}_{\mathrm{CM}} \neq 0.$$

Proof. Sketch of Proof. Using (A47), proportionality would require $\rho_X \equiv \mu^{(N)}_X / \mu^{(\nabla N)}_X$ 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 $X = \mathrm{ROD}, \mathrm{CM}$ unless $c^{(Y)}_{X,\mathcal{H}} = c^{(Y)}_{X,\mathcal{Y}} = 0$ (static trivial branch $\mathcal{H} = \mathcal{Y} = 0$, $\nabla\hat{n} = 0$). Hence, $\boldsymbol{r}_1 \nparallel \boldsymbol{r}_2$ generically. □

Small-k structure (symbolic).

On weak-field patches, one finds

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

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)

$$J^\mu_\Delta = \frac{M_{\rm Pl}^2}{8}\begin{pmatrix} J^0_\Delta = \;\; \Pi^{ij}\, h^{\rm TT}_{ij} \\ J^k_\Delta = -\,\Xi^{ij}\, \partial_k h^{\rm TT}_{ij} \end{pmatrix}, \tag{A52}$$

with $\Pi^{ij} \equiv \delta(\delta^2\mathcal{L})/\delta\dot{h}^{\rm TT}_{ij}$ and $\Xi^{ij} \equiv \delta(\delta^2\mathcal{L})/\delta(\partial^2 h^{\rm TT}_{ij})$. Using TT conditions, the rank-one traceless normalization, and the bulk equality of the routes (Section 5),

$$\boxed{K(w) - G(w) = \frac{1}{a^3}\,\partial_\mu(a^3 J^\mu_\Delta).} \tag{A53}$$

Under $\delta_\xi h^{\rm TT}_{ij} = 0$, $\delta_\xi h_{0\mu} = \partial_\mu \xi_0 + \cdots$, the representative shifts as $\delta_\xi J^\mu_\Delta = \nabla_\nu X^{[\mu\nu]}[\xi]$; 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),

$$J^0_\Delta = \;\frac{A_\star\,\eta^2}{2}\,\Sigma_\epsilon\, h^{\rm TT}_{ij}\,\dot{h}^{\rm TT}_{ij} \;+\; \nabla_i(\ldots), \tag{A54}$$

$$J^k_\Delta = -\,\frac{A_\star\,\eta^2}{2}\,\Sigma_\epsilon\,\frac{1}{a^2}\,h^{\rm TT}_{ij}\,\partial_k h^{\rm TT}_{ij} \;+\; \partial_\ell(\ldots)^{[k\ell]}, \tag{A55}$$

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

$$J^0_\Delta = \frac{A_\star\,\eta^2}{2}\,\Sigma_\epsilon\, h^{\rm TT}_{ij}\,\dot{h}^{\rm TT}_{ij} + \partial_i(\ldots)^i, \quad J^k_\Delta = -\frac{A_\star\,\eta^2}{2}\,\Sigma_\epsilon\, h^{\rm TT}_{ij}\,\partial_k h^{\rm TT}_{ij} + \partial_\ell(\ldots)^{[k\ell]}.$$

Boundary consequence and luminality.

With the variational domain above (A4), $\int a^3(K - G) = \int \partial_\mu(a^3 J^\mu_\Delta) = 0$. At the *locked* weights $w^\star$ (Section 6; 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 $w^\star_{\rm ROD}/w^\star_{\rm CM}$ (up to the GR normalization). At these weights, $K - G = a^{-3}\partial_\mu(a^3 J^\mu_\Delta)$ implies exact luminality.

Appendix D (at a glance).

- *Mixing matrix.* Four dimensionless coefficients $\mu^{(Y)}_X$ ($X = {\rm ROD}, {\rm CM}$; $Y = N, \nabla N$) 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 ($\Delta \sim \alpha_0\eta^4\Sigma_\epsilon^2 k^2 + \cdots$).
- *Equal-coefficient identity.* $K - G = a^{-3}\partial_\mu(a^3 J^\mu_\Delta)$; under A4, this yields $K = G$ and $c_T^2 = 1$ at the locked weights $w^\star$.

Reproducibility Note

Tables A1 and 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, and (iii) projects onto the bilinears of Equations (A43) and (A44). Scripts and hashes that also produced Figures 7 and 8 are cataloged in the repository note described in the Data Availability Statement. (Code: `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 $\eta_{AB} = \mathrm{diag}(-,+,+,+)$, and adopt the gamma-matrix conventions

$$\{\gamma^A, \gamma^B\} = 2\eta^{AB}, \qquad \gamma_5 \equiv i\,\gamma^0\gamma^1\gamma^2\gamma^3, \qquad \epsilon^{0123} = +1, \tag{A56}$$

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

$$\mathcal{L}_\psi = e\,\bar\psi\Big(i\,e_A{}^\mu\gamma^A D_\mu - m\Big)\psi, \qquad D_\mu\psi = \partial_\mu\psi + \tfrac{1}{4}\,\omega^{AB}{}_\mu\gamma_{AB}\psi, \qquad e \equiv \sqrt{-g}. \tag{A57}$$

Splitting the spin connection into Levi–Civita plus contorsion, $\omega^{AB}{}_\mu = \omega^{AB}{}_\mu(\mathrm{LC}) + K^{AB}{}_\mu$, the *torsion-dependent* piece of (A57) is

$$\Delta\mathcal{L}_\psi[K] = \frac{i}{4}\,e\,\bar\psi\,e_A{}^\mu\gamma^A\,K^{BC}{}_\mu\gamma_{BC}\,\psi = \frac{i}{4}\,e\,\bar\psi\,\gamma^A\,K_{BCA}\,\gamma^{BC}\,\psi, \tag{A58}$$

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}):

$$K_{ABC} = \frac{1}{3}(\eta_{AC}T_B - \eta_{AB}T_C) - \frac{1}{6}\,\epsilon_{ABCD}\,S^D + q_{ABC}, \qquad T_A \equiv T^B{}_{AB}, \quad S^A \equiv \epsilon^{ABCD}T_{BCD}, \tag{A59}$$

with $q^B{}_{AB} = 0$ and $\epsilon^{ABCD}q_{BCD} = 0$. Using the gamma identities listed below Equation (A56) and the antisymmetry of K_{BCA} in (B, C), Equation (A58) reduces to

$$\Delta\mathcal{L}_\psi[K] = e\Big(c_A\,S_\mu J_5^\mu + c_V\,T_\mu J^\mu\Big) \;+\; e\,\mathcal{C}^{\mu\nu\rho}(q)\;\bar\psi\,\gamma_{[\mu}\gamma_\nu\gamma_{\rho]}\,\psi, \tag{A60}$$

where $J^\mu \equiv \bar\psi\gamma^\mu\psi$ and $J_5^\mu \equiv \bar\psi\gamma^\mu\gamma_5\psi$. In our normalization (A56), the *numerical* coefficients are

$$\boxed{c_A = \frac{3}{8}, \qquad c_V = -\frac{1}{4}} \tag{A61}$$

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

$$\boxed{\Delta\mathcal{L}_\psi[T, S] = e\Big(\tfrac{3}{8}\,S_\mu J_5^\mu - \tfrac{1}{4}\,T_\mu J^\mu\Big),} \tag{A62}$$

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

Cross-check.

Equation (A62) 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 (A61) 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 (16), the axial and traceless torsion irreps vanish, $S^\mu = 0, \quad q_{\lambda\mu\nu} = 0$, and the trace aligns with the spurion gradient, $T_\mu = 3\eta\,\partial_\mu\epsilon$. Equation (A62) therefore reduces to the single-trace channel:

$$\mathcal{L}^{(\mathrm{trace})}_{\psi,T} = e\,c_V\,T_\mu J^\mu = 3c_V\eta\,e\,(\partial_\mu\epsilon)\,J^\mu. \tag{A63}$$

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

$$\mathcal{L}_\psi[\psi',\bar\psi'] = \mathcal{L}_\psi[\psi,\bar\psi] - \alpha\,e\,(\partial_\mu\epsilon)\,J^\mu \ + \ \nabla_\mu(\cdots), \tag{A64}$$

so choosing $\alpha = 3c_V\eta\ (= -\frac{3}{4}\eta$ *in our normalization*$)$ *cancels* (A63) *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.*

Proof. Sketch of Proof. Insert $\psi' = e^{i\alpha\epsilon}\psi$ into the kinetic term $e\,\bar\psi i e_A{}^\mu\gamma^A\partial_\mu\psi$ and use the Leibniz rule. The derivative acting on $e^{i\alpha\epsilon}$ generates the shift $-\alpha e(\partial_\mu\epsilon)J^\mu$ in (A64); 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). □

Combining (A63) with Proposition A5 and the coefficient assignment (A61) 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 Proposition 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 Sections 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 Proposition 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 (A64) 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	$\Delta\mathcal{L}_\psi = e\left(\frac{3}{8}S_\mu J_5^\mu - \frac{1}{4}T_\mu J^\mu\right)$
C1 (pure trace)	$S^\mu = 0$, $q_{\lambda\mu\nu} = 0$, $T_\mu = 3\eta\,\partial_\mu\epsilon$
Vector rephasing	$\psi \to e^{i\alpha\epsilon}\psi$ 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 $\psi \to e^{i(3c_V\eta)\epsilon}\psi$, 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 (Theorem 1; Section 2.6). (Chern–Simons modified gravity is the canonical benchmark for helicity-split propagation; see, e.g., Jackiw–Pi; Alexander–Yunes.)

- **(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 (Theorem 2); Figures 2 and 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 $A_\star\sqrt{-g}\ \mathrm{sgn}(\Sigma_\epsilon)\ I_T$ 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 $K-G = \partial_\mu(a^{-3}J^\mu_\Delta a^3)$ 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.
- **(D7) NLO slope.** At next order, the leading PT-even dispersion predicts $\delta c_T^2(k) = b\,k^2/\Lambda^2$ (Equation (43)): a log–log slope $\simeq 2$ in clean frequency windows is characteristic (Figure 11; Section 8.2). (Caveat: finite-window fits may be biased by instrumental systematics or astrophysical priors; interpret bounds conservatively.)
- **(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 8.5):
  - **(D9-a) Quadratic scaling.** Fit $|\delta c_T^2(k)| \propto k^2$ on the admissible band ($k \ll \Lambda$) using Equation (45). 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 $\Delta A_\star(k) \equiv A_\star^{(rank-onedeterminantroute)} - A_\star^{(\mathrm{CM})} \to 0$ and boundary flux-ratio unity $\mathcal{R}_{rank-onedeterminantroute/\mathrm{CM}} \to 1$ 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 Appendix F.3 for caveats and representative citations.

Table A3 summarizes the operational disambiguation diagnostics used in Appendix F.2.

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

Table A3. *Cont.*

Diagnostic	This Work	CS-Mod. Grav.	Horndeski /DHOST	Teleparallel $/f(T)$	EC/ ECSK-like
c_T at LO	=1 by identity (D5)	Helicity-split	Tuned $\to 1$	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 $\mathcal{R}_{X/Y}$	$\to 1$ (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

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 Theorem 1. Representative anchors can be found in [32,33].

Horndeski/DHOST.

Post-GW170817 constraints force $c_T \to 1$ by parameter tuning or by restricting to subclasses; however, no equal-coefficient *identity* of the type $K-G = \partial_\mu(\cdots)$ is generally present, and an extra scalar DoF is typical. See the multimessenger bounds and reviews [4–8,36].

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 [12–14,37].

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 include covariant phase space and surface charges [26,27,38], and the Holst/Nieh–Yan parity bookkeeping [21,22,24,25].

Representative anchors (one-line map).

Metric-affine/Einstein–Cartan: [12–14]. Holst/Nieh–Yan: [21,22,24,25]. Covariant phase space & boundary charges: [26,27,38]. CS-modified gravity: [32,33]. GW170817 speed bounds & implications: [1–6,36,39].

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

$$\boxed{\mathcal{T}_\mu \;\equiv\; T_\mu - \partial_\mu \epsilon, \qquad \Gamma^\alpha{}_{\mu\nu} \to \Gamma^\alpha{}_{\mu\nu} + \delta^\alpha_\mu \xi_\nu : \; \mathcal{T}_\mu \text{ invariant.}} \tag{A65}$$

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. Stueckelberg Completion and the Spurion Limit

A projectively invariant, two-derivative completion is

$$\boxed{S_{\rm Stk} = \int d^4x \, \sqrt{-g}\Big[\tfrac{M_P^2}{2} R(e,\omega) - \tfrac{m_T^2}{2} \mathcal{T}_\mu \mathcal{T}^\mu - \tfrac{f_\epsilon^2}{2} (\partial\epsilon)^2\Big], \qquad \mathcal{T}_\mu \equiv T_\mu - \partial_\mu \epsilon,} \tag{A66}$$

with $m_T > 0$ and $f_\epsilon \geq 0$ real. The mass-like 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 $\omega = \mathring{\omega} + K(T)$ and using the standard algebraic split of $R(e,\omega)$ into $R(\mathring{\omega})$ 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 (A66),

$$\delta_\omega S_{\rm Stk} \;\Rightarrow\; S^\mu = 0, \qquad q_{\lambda\mu\nu} = 0, \qquad m_T^2 \, (T_\mu - \partial_\mu \epsilon) + \frac{4\alpha_1}{3} \, T_\mu = 0, \tag{A67}$$

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

$$\boxed{T_\mu = \partial_\mu \epsilon, \qquad S^\mu = 0, \qquad q_{\lambda\mu\nu} = 0,} \tag{A68}$$

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

Two equivalent implementations of the *spurion limit*.

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

$$\boxed{\begin{aligned} &\text{(i) Hard penalty: } m_T \to \infty \;\Rightarrow\; \mathcal{T}_\mu = 0, \\ &\text{(ii) Lagrange current: } \mathcal{L}_\Lambda = \sqrt{-g}\, \Lambda^\mu \mathcal{T}_\mu. \end{aligned}} \tag{A69}$$

In case (i), integrating 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 $f_\epsilon^2 \to 0^+$ (or treating ϵ as non-dynamical) *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,

$$S_{\rm eff}^{(\rm LO)} = \int d^4x\sqrt{-g}\Big[\tfrac{M_P^2}{2}R(\mathring{\omega}) + \alpha_1\, I_T + \alpha_2\,\Sigma_\epsilon\Big] \;+\; \mathcal{O}\!\left(\frac{\partial^4}{m_T^2}\right), \qquad I_T \equiv -\tfrac{1}{4}\,\Pi_{\rm PT}[T^A{}_{BC}T_A{}^{BC}], \tag{A70}$$

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 (A66) by a topological coupling

$$S_{\text{t-ax}} = \int d^4x\sqrt{-g}\,\frac{\alpha}{4}\,\epsilon\,\mathrm{NY}, \qquad \mathrm{NY} \equiv d\Big(e^A \wedge T_A\Big), \tag{A71}$$

which is exact. Under A5, NY is a *boundary convention*; integrating by parts and using (A68) gives bulk terms proportional to $\partial_\mu\epsilon$ contracted with improvement currents and, 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).

(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 $\omega_\pm = \mathring{\omega} \pm \frac{1}{2}\,H$. 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 $\epsilon \to \epsilon + c$ protects the *gradient-only* appearance of ϵ and suppresses f_ϵ^2 technically: taking $f_\epsilon^2 \to 0$ *enhances* symmetry (exact spurion posture).

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

https://doi.org/10.3390/sym18010170

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 $\psi \to e^{i\alpha\epsilon}\psi$ (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

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

algebraically enforces the pure-trace alignment $T_\mu = \partial_\mu\epsilon$ (C1) and $S^\mu = q_{\lambda\mu\nu} = 0$; see Equations (A67)–(A68). Taking the *spurion limit* (hard penalty $m_T \to \infty$ or a Lagrange current for $\mathcal{T}_\mu$, with $f_\epsilon^2 \to 0$ at two derivatives) yields precisely the observable sector used in the main text: PT-even scalars depend only on $\mathcal{T}_\mu$ (and 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 Section 9. The statement used in Section 9—"Treating ϵ as non-dynamical is the low-energy limit of a Stueckelberg completion in which $m_T \to \infty$ freezes $\mathcal{T}_\mu$; complementary EFT constructions reduce to the same gradient-only dependence"—is precisely the content of Equations (A66)–(A69).

Appendix H. Four-Derivative Basis: Operator-by-Operator Tensor Reduction

This appendix provides the auditable reduction behind Table 2. We expand each PT-even, projectively consistent four-derivative scalar density to quadratic order in transverse-traceless (TT) tensor perturbations and record its footprint on the quadratic tensor action. The outcome is a basis-level statement: within the symmetry filter and admissible-domain posture (A4–A5), all operators either (i) are boundary terms, (ii) are TT-neutral at quadratic order, or (iii) renormalize the TT kinetic and gradient terms in lockstep; only one direction yields a differential deformation that is captured by the NLO parameter b/Λ^2 in Section 8.2.

Appendix H.1. Tensor Setup and Coefficient Bookkeeping

We consider a smooth background metric $\bar{g}_{\mu\nu}$ and TT fluctuations h_{ij} in ADM language. For concreteness (and because it is the data-facing case), one may take a spatially flat FRW patch,

$$ds^2 = -dt^2 + a^2(t)\left(\delta_{ij} + h_{ij}\right)dx^i dx^j, \qquad \partial_i h_{ij} = 0, \quad h_{ii} = 0, \tag{A72}$$

though the classification below is local and extends to any smooth background where the TT decomposition is well-defined.

The quadratic tensor action is written as

$$S_T^{(2)} = \frac{M_T^2}{8} \int dt\, d^3x\, a^3 \Big[K\, \dot{h}_{ij}\dot{h}_{ij} - G\, \frac{\partial_k h_{ij}\partial_k h_{ij}}{a^2} + \cdots \Big], \tag{A73}$$

where M_T^2 is the tensor normalization (equal to $M_{\rm Pl}^2$ in GR) and the ellipsis denotes slow-roll/background-suppressed mass-like terms that do not affect the high-frequency propagation speed. The tensor speed is

$$c_T^2 = \frac{G}{K}. \tag{A74}$$

At LO (two-derivative order), the locking mechanism of Section 6 yields $K = G$ and thus $c_T = 1$.

At NLO, we add the four-derivative correction

$$\Delta S^{(4)} \;=\; \frac{1}{\Lambda^2} \int d^4x\, \sqrt{-g} \sum_i c_i\, \mathcal{O}_i^{(4)}, \tag{A75}$$

and expand each $\mathcal{O}_i^{(4)}$ to quadratic order in h_{ij}. For each operator, we record the induced shifts

$$K \to K + \Delta K_i, \qquad G \to G + \Delta G_i, \tag{A76}$$

and whether an additional $(\nabla^2 h)^2$-type term appears, which yields the slope-2 dispersion of Section 8.2.

Appendix H.2. Operator-by-Operator Reduction

(i) Gauss–Bonnet: a boundary density in 4D

The Gauss–Bonnet combination

$$\mathcal{O}_{\rm GB}^{(4)} = R_{\mu\nu\rho\sigma}R^{\mu\nu\rho\sigma} - 4R_{\mu\nu}R^{\mu\nu} + R^2 \tag{A77}$$

is a total derivative in $d = 4$,

$$\sqrt{-g}\, \mathcal{O}_{\rm GB}^{(4)} = \partial_\mu \big(\sqrt{-g}\, J_{\rm GB}^\mu \big), \tag{A78}$$

so on admissible domains (A4–A5), it contributes only a boundary flux and does not enter the bulk quadratic tensor action. Hence,

$$\Delta K_{\rm GB} = \Delta G_{\rm GB} = 0, \qquad \text{(bulk TT, quadratic order)}. \tag{A79}$$

(ii) Pure curvature squares: R^2 and $R_{\mu\nu}R^{\mu\nu}$

For TT perturbations, the linearized scalar curvature satisfies $\delta R = 0$ (it depends only on the trace/divergence components). As a consequence, the quadratic expansion of R^2 about a smooth background reduces to a background-renormalization term,

$$\delta^{(2)}\big(\sqrt{-g}\, R^2\big) = 2\sqrt{-\bar{g}}\, \bar{R}\, \delta^{(2)}R \;+\; \text{(TT boundary/TT-neutral)}, \tag{A80}$$

and similarly for $R_{\mu\nu}R^{\mu\nu}$. In the EFT posture of this paper, these contributions can be absorbed into (i) a renormalization of background equations and (ii) a lockstep renormalization of the TT kinetic and gradient coefficients, without generating a differential $K \neq G$. Concretely, one finds

$$\Delta S_{R^2}^{(2)} + \Delta S_{RR}^{(2)} = \frac{M_T^2}{8} \int dt\, d^3x\, a^3\, \zeta(t) \Big[\dot{h}_{ij}\dot{h}_{ij} - \frac{\partial_k h_{ij}\partial_k h_{ij}}{a^2} \Big] + \text{(TT boundary/TT mass-like)}, \tag{A81}$$

for some background-dependent $\zeta(t) = \mathcal{O}(c_i\,\bar{R}/\Lambda^2)$. Therefore,

$$\Delta K_{R^2} = \Delta G_{R^2}, \qquad \Delta K_{RR} = \Delta G_{RR}, \tag{A82}$$

and these operators do not generate a leading NLO deviation in c_T.

(iii) Spurion-dressed curvature: $\Sigma_\epsilon\, R$

The operator

$$\mathcal{O}^{(4)}_{\Sigma_\epsilon R} = \Sigma_\epsilon\, R \tag{A83}$$

acts as a background-dependent rescaling of the Einstein–Hilbert term at quadratic order in TT, because Σ_ϵ is a scalar background in the spurion posture (A6). Expanding to quadratic order yields

$$\Delta S^{(2)}_{\Sigma_\epsilon R} = \frac{M_T^2}{8}\int dt\, d^3x\, a^3\, \chi(t)\Big[\dot{h}_{ij}\dot{h}_{ij} - \frac{\partial_k h_{ij}\partial_k h_{ij}}{a^2}\Big] + (\text{TT boundary/TT mass-like}), \tag{A84}$$

with $\chi(t) = \mathcal{O}(c_{\Sigma_\epsilon R}\Sigma_\epsilon/\Lambda^2)$. Hence,

$$\Delta K_{\Sigma_\epsilon R} = \Delta G_{\Sigma_\epsilon R}, \tag{A85}$$

and $\Sigma_\epsilon R$ preserves $c_T = 1$ at this order.

(iv) Pure spurion operators: Σ_ϵ^2 and $(\nabla\Sigma_\epsilon)^2$

The operators

$$\mathcal{O}^{(4)}_{\Sigma_\epsilon^2} = \Sigma_\epsilon^2, \qquad \mathcal{O}^{(4)}_{\nabla\Sigma_\epsilon} = (\nabla_\mu\Sigma_\epsilon)(\nabla^\mu\Sigma_\epsilon) \tag{A86}$$

contain no metric derivatives beyond those implicit in contractions. In the strict spurion limit (A6), they contribute only to the scalar/background sector and are TT-neutral at quadratic order about a fixed background:

$$\Delta K_{\Sigma_\epsilon^2} = \Delta G_{\Sigma_\epsilon^2} = 0, \qquad \Delta K_{\nabla\Sigma_\epsilon} = \Delta G_{\nabla\Sigma_\epsilon} = 0 \qquad \text{(fixed background, TT quadratic order)}. \tag{A87}$$

If ϵ is promoted to a dynamical field, these operators govern the scalar channel and can affect tensors only through mixing operators discussed in Section 8.4; this lies beyond the strict posture.

(v) The unique differential direction: $C_{\mu\nu\rho\sigma}C^{\mu\nu\rho\sigma}$

Finally, consider

$$\mathcal{O}^{(4)}_{C^2} = C_{\mu\nu\rho\sigma}C^{\mu\nu\rho\sigma}, \tag{A88}$$

with $C_{\mu\nu\rho\sigma}$ the Weyl tensor of the Levi–Civita connection. About conformally flat backgrounds (including FRW), the background Weyl tensor vanishes, so the leading contribution to the quadratic action is simply the square of the linearized Weyl tensor. In TT gauge, this yields a genuine $(\nabla^2 h)^2$-type correction in the reduced quadratic action,

$$\Delta S^{(2)}_{C^2} = \frac{M_T^2}{8}\int dt\, d^3x\, a^3\,\Big[\tilde{K}\,\dot{h}_{ij}\dot{h}_{ij} - \tilde{G}\,\frac{\partial_k h_{ij}\partial_k h_{ij}}{a^2} + \tilde{B}\,\frac{(\partial^2 h_{ij})(\partial^2 h_{ij})}{a^4}\Big], \tag{A89}$$

where $\tilde{K}, \tilde{G} = \mathcal{O}(H^2/\Lambda^2)$ are background-suppressed and renormalize in lockstep at leading propagation order, while the last term is the unsuppressed four-spatial-derivative correction with

$$\tilde{B} = \mathcal{O}\Big(\frac{c_{C^2}}{\Lambda^2}\Big). \tag{A90}$$

In Fourier space, the $(\partial^2 h)^2$ term contributes $\propto k^4 h^2$ and therefore deforms the dispersion relation as

$$\omega^2 = k^2_{\rm phys}\Big[1 + b\,\frac{k^2_{\rm phys}}{\Lambda^2}\Big], \qquad b \;\propto\; c_{C^2}, \tag{A91}$$

which is the slope-2 target used in Section 8.2. Crucially, this operator is the only member of the complete PT-even, projectively consistent four-derivative basis that generates the $(\partial^2 h)^2$ structure; all other operators are boundary/TT-neutral/lockstep as shown above. This is the basis-level reason Table 2 singles out $\mathcal{O}^{(4)}_{C^2}$ as the unique "dangerous" direction.

Appendix H.3. Summary: Footprint Table in $(\Delta K, \Delta G, \Delta B)$ *Form*

Collecting the results, the quadratic TT footprints can be summarized as

$$\Delta S^{(2)}_T = \frac{M_T^2}{8}\int dt\, d^3x\, a^3\Big[(\Delta K)\,\dot h^2_{ij} - (\Delta G)\,\frac{(\partial h_{ij})^2}{a^2} + (\Delta B)\,\frac{(\partial^2 h_{ij})^2}{a^4}\Big], \tag{A92}$$

with the operator-by-operator entries:

$$\begin{aligned}
&\mathcal{O}^{(4)}_{\rm GB}: \quad \Delta K = \Delta G = \Delta B = 0 \quad \text{(bulk, A4–A5)},\\
&\mathcal{O}^{(4)}_{R^2},\ \mathcal{O}^{(4)}_{RR},\ \mathcal{O}^{(4)}_{\Sigma_\epsilon R}: \quad \Delta K = \Delta G,\ \Delta B = 0,\\
&\mathcal{O}^{(4)}_{\Sigma^2_\epsilon},\ \mathcal{O}^{(4)}_{\nabla\Sigma_\epsilon}: \quad \Delta K = \Delta G = \Delta B = 0 \quad \text{(fixed background, spurion posture)},\\
&\mathcal{O}^{(4)}_{C^2}: \quad \Delta B \neq 0 \;\Rightarrow\; \delta c_T^2(k) = b\,k^2_{\rm phys}/\Lambda^2 \text{ (slope-2)}.
\end{aligned} \tag{A93}$$

This completes the auditable basis argument required for the NLO signature and for the technical-naturalness discussion: within the symmetry-selected four-derivative basis, any differential deformation of tensor propagation is funneled into a single operator direction.

References

1. Abbott, B.P.; Abbott, R.; Abbott, T.D.; Acernese, F.; Ackley, K.; Adams, C.; Adams, T.; Addesso, P.; Adhikari, R.X.; Adya, V.B.; et al. GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral. *Phys. Rev. Lett.* **2017**, *119*, 161101. [CrossRef]
2. Goldstein, A.; Veres, P.; Burns, E.; Briggs, M.S.; Hamburg, R.; Kocevski, D.; Wilson-Hodge, C.A.; Preece, R.D.; Poolakkil, S.; Roberts, O.J.; et al. An Ordinary Short Gamma-Ray Burst with Extraordinary Implications: Fermi-GBM Detection of GRB 170817A. *Astrophys. J. Lett.* **2017**, *848*, L14.
3. Abbott, B.P.; Abbott, R.; Abbott, T.D.; Acernese, F.; Ackley, K.; Adams, C.; Adams, T.; Addesso, P.; Adhikari, R.X.; Adya, V.B.; et al. Multi-messenger Observations of a Binary Neutron Star Merger. *Astrophys. J. Lett.* **2017**, *848*, L12. [CrossRef]
4. Creminelli, P.; Vernizzi, F. Dark Energy after GW170817 and GRB170817A. *Phys. Rev. Lett.* **2017**, *119*, 251302. [CrossRef]
5. Baker, T.; Bellini, E.; Ferreira, P.G.; Lagos, M.; Noller, J.; Sawicki, I. Strong constraints on cosmological gravity from GW170817 and GRB170817A. *Phys. Rev. Lett.* **2017**, *119*, 251301. [CrossRef] [PubMed]
6. Ezquiaga, J.M.; Zumalacárregui, M. Dark Energy After GW170817: Dead Ends and the Road Ahead. *Phys. Rev. Lett.* **2017**, *119*, 251304. [CrossRef]
7. Kase, R.; Tsujikawa, S. Dark energy in Horndeski theories after GW170817: A review. *Int. J. Mod. Phys. D* **2019**, *28*, 1942005. [CrossRef]
8. Langlois, D. Dark Energy and Modified Gravity in Degenerate Higher-Order Scalar–Tensor (DHOST) theories. *Int. J. Mod. Phys. D* **2019**, *28*, 1942006. [CrossRef]
9. Clifton, T.; Ferreira, P.G.; Padilla, A.; Skordis, C. Modified gravity and cosmology. *Phys. Rep.* **2012**, *513*, 1–189. [CrossRef]
10. Heisenberg, L. A systematic approach to generalisations of General Relativity and their cosmological implications. *Phys. Rep.* **2019**, *796*, 1–113. [CrossRef]
11. Palatini, A. Deduzione invariantiva delle equazioni gravitazionali dal principio di Hamilton. *Rend. Circ. Mat. Palermo* **1919**, *43*, 203–212. [CrossRef]
12. Hehl, F.W.; von der Heyde, P.; Kerlick, G.D.; Nester, J.M. General Relativity with Spin and Torsion: Foundations and Prospects. *Rev. Mod. Phys.* **1976**, *48*, 393–416. [CrossRef]

13. Hehl, F.W.; McCrea, J.D.; Mielke, E.W.; Ne'eman, Y. Metric-affine gauge theory of gravity: Field equations, Noether identities, world spinors, and breaking of dilaton invariance. *Phys. Rep.* **1995**, *258*, 1–171. [CrossRef]
14. Shapiro, I.L. Physical aspects of the space–time torsion. *Phys. Rep.* **2002**, *357*, 113–213. [CrossRef]
15. Kibble, T.W.B. Lorentz invariance and the gravitational field. *J. Math. Phys.* **1961**, *2*, 212–221. [CrossRef]
16. Sciama, D.W. On the analogy between charge and spin in general relativity. In *Recent Developments in General Relativity*; Infeld, L., Ed.; Pergamon Press: Oxford, UK; PWN: Warsaw, Poland, 1962; pp. 415–439.
17. Hammond, R.T. Torsion gravity. *Rep. Prog. Phys.* **2002**, *65*, 599–649. [CrossRef]
18. Bahamonde, S.; Dialektopoulos, K.F.; Said, C.L.; Said, J.L. Teleparallel gravity: From theory to cosmology. *Rep. Prog. Phys.* **2023**, *86*, 026901. [CrossRef] [PubMed]
19. Cai, Y.-F.; Capozziello, S.; Laurentis, M.D.; Saridakis, E.N. $f(T)$ teleparallel gravity and cosmology. *Rep. Prog. Phys.* **2016**, *79*, 106901. [CrossRef]
20. Chen, C.-C. PT-Projected Projective Palatini Gravity: Two-Derivative Operator Basis, Admissible Equivalences, and a Local IR Residual. *Preprints* **2026**, *2026*, 2026010352. [CrossRef]
21. Nieh, H.T.; Yan, M.L. An identity in Riemann–Cartan geometry. *J. Math. Phys.* **1982**, *23*, 373–374. [CrossRef]
22. Chandia, O.; Zanelli, J. Topological invariants, instantons and the chiral anomaly on spaces with torsion. *Phys. Rev. D* **1997**, *55*, 7580–7585. [CrossRef]
23. Arnowitt, R.; Deser, S.; Misner, C.W. The Dynamics of General Relativity. In *Gravitation: An Introduction to Current Research*; Witten, L., Ed.; Wiley: New York, NY, USA, 1962; pp. 227–265; reprinted in *Gen. Rel. Grav.* **2008**, *40*, 1997–2027.
24. Holst, S. Barbero's Hamiltonian derived from a generalized Hilbert–Palatini action. *Phys. Rev. D* **1996**, *53*, 5966–5969. [CrossRef]
25. Mercuri, S. Fermions in the Ashtekar–Barbero connection formalism: The Nieh–Yan invariant as a source of the Immirzi parameter. *Phys. Rev. D* **2006**, *73*, 084016. [CrossRef]
26. Iyer, V.; Wald, R.M. Some properties of Noether charge and a proposal for dynamical black hole entropy. *Phys. Rev. D* **1994**, *50*, 846–864. [CrossRef] [PubMed]
27. Regge, T.; Teitelboim, C. Role of surface integrals in the Hamiltonian formulation of general relativity. *Ann. Phys.* **1974**, *88*, 286–318. [CrossRef]
28. Olmo, G.J. Palatini Approach to Modified Gravity: $f(R)$ Theories and Beyond. *Int. J. Mod. Phys. D* **2011**, *20*, 413–462. [CrossRef]
29. Flanagan, E.E. Palatini form of $1/R$ gravity. *Phys. Rev. Lett.* **2004**, *92*, 071101. [CrossRef]
30. Barausse, E.; Sotiriou, T.P.; Miller, J.C. Curvature singularities in Palatini $f(R)$ gravity. *Phys. Rev. D* **2008**, *77*, 104035.
31. Sotiriou, T.P.; Faraoni, V. $f(R)$ theories of gravity. *Rev. Mod. Phys.* **2010**, *82*, 451–497. [CrossRef]
32. Jackiw, R.; Pi, S.-Y. Chern–Simons modification of general relativity. *Phys. Rev. D* **2003**, *68*, 104012. [CrossRef]
33. Alexander, S.; Yunes, N. Chern–Simons Modified Gravity. *Phys. Rep.* **2009**, *480*, 1–55. [CrossRef]
34. Obukhov, Y.N. Gravitational waves in Poincaré gauge gravity theory. *Phys. Rev. D* **2017**, *95*, 084028. [CrossRef]
35. Elizalde, E.; Odintsov, S.D.; Obukhov, V.V. Gravitational waves in Einstein–Cartan theory. *Phys. Dark Univ.* **2023**, *41*, 101256.
36. Sakstein, J.; Jain, B. Implications of the Neutron Star Merger GW170817 for Cosmological Scalar–Tensor Theories. *Phys. Rev. Lett.* **2017**, *119*, 251303. [CrossRef] [PubMed]
37. Hehl, F.W.; Datta, B.K. Nonlinear spinor equation and asymmetric connection in general relativity. *J. Math. Phys.* **1971**, *12*, 1334–1339. [CrossRef]
38. Wald, R.M. *General Relativity*; University of Chicago Press: Chicago, IL, USA, 1984.
39. Savchenko, V.; Ferrigno, C.; Kuulkers, E.; Bazzano, A.N.; Bozzo, E.; Brandt, S.; Chenevez, J.; Courvoisier, T.L.; Diehl, R.; Domingo, A.; et al. INTEGRAL Detection of the First Prompt Gamma-Ray Signal Coincident with the Gravitational-wave Event GW170817. *Astrophys. J. Lett.* **2017**, *848*, L15. [CrossRef]