

A Structural No-Go Survey of Minimal Spin-Torsion Routes to Dark Energy and Bounce Phenomenology

Houston Golden^{1,*}

¹*Independent Researcher, Los Angeles, California, USA*

(Dated: August 6, 2026 (v1C.0.7))

We present a structural no-go survey of the minimal Einstein–Cartan–Holst (ECH) spin-torsion framework as a source of late-time dark energy and bounce phenomenology. Working from the algebraic torsion-elimination setup established in a companion paper, we catalog thirteen distinct mechanism-class constraints — fourteen historical catalog entries, one subsumed by another — spanning seven foundational mechanism classes and six observational-channel branches, each closing one or more of the four routes by which the ECH bounce could plausibly source a Λ -like late-time density. Of the four parity-odd/dark-energy channels enumerated by this framework, we present amplitude-budget closures of the two historical quantum routes: Route 2 (one-loop graviton-sector corrections to the Holst term), whose budget is anchored in the published one-loop renormalization of the Holst-plus-fermion sector and closes against the *observed birefringence amplitude* with roughly sixty orders of magnitude (conservatively ≥ 58) of suppression margin, and Route 3 (quantum running of the Barbero–Immirzi parameter), bounded by an integrated one-loop renormalization-group flow that leaves the contribution 61–67 orders of magnitude below the *observed dark-energy density* (the ~ 67 -order endpoint from the derived integrated flow, the ~ 61 -order endpoint from a deliberately pessimistic chiral-count bound). Neither route is presented as a viable mapping: the companion paper does not retain these historical mappings as results, and this survey closes them as bounded amplitude budgets. We further present an operator-basis argument: a six-member basis $\{O1\text{--}O6\}$ of local, gauge-invariant, diffeomorphism-covariant parity-odd densities of mass dimension exactly four, exhibited under stated construction rules for the minimal-coupling ECH field content, is shown to close — every member is a topological total derivative, a Fierz-closed four-fermion contact term uniformly suppressed by M_{Pl}^{-2} , or identically vanishing by the algebraic Bianchi identity; completeness of the basis is asserted from the construction rules, not proved by exhaustive symbolic enumeration. The catalog carries an explicit evidentiary classification (rigorous result / structural argument / ansatz-level estimate): only the perturbation-transparency result is a Tier-I rigorous theorem, and the survey is a channel-level, not operator-level, closure within a stated minimal-coupling scope — non-minimal completions, an added light scale or non-minimal torsion coupling, lie explicitly outside it.

I. INTRODUCTION

Minimal Einstein–Cartan–Holst (ECH) gravity is among the simplest torsionful extensions of general relativity motivated by loop quantum cosmology: promoting the spin connection to an independent field, with torsion sourced algebraically by fermion spin currents, introduces exactly one new parameter (the Barbero–Immirzi parameter γ) and one new coupling scale ($\kappa = 8\pi G \sim M_{\text{Pl}}^{-2}$), with no new propagating degree of freedom. In loop-quantized cosmology the resulting bounce replaces the classical singularity with a quantum bounce at Planck-scale density. A recurring question in this framework, and in torsion-based bounce cosmology more generally, is whether the same minimal spin-torsion sector that produces the bounce can also source the observed late-time dark-energy density $\rho_{\Lambda,\text{obs}} \approx (2.25 \text{ meV})^4$.

This paper answers that question systematically. We enumerate every distinct mechanism class by which minimal ECH could plausibly connect its bounce-scale torsion sector to a late-time Λ -like density — seven structural

“foundations” (parameter-space, symmetry, and naturalness arguments) and six observational branches (CMB, gravitational-wave, and relic-abundance channels) — and show that each is closed by an explicit, individually labeled argument: an amplitude-suppression bound, a naturalness/fine-tuning objection, a decoupling identity, or a classification argument. The resulting fourteen-entry catalog collapses to thirteen distinct mechanism-class constraints (one entry is subsumed by a more complete result) and is presented alongside two amplitude-budget closures — Route 2 (one-loop graviton-sector corrections to the Holst term) and Route 3 (quantum running of the Barbero–Immirzi parameter) — together with an operator-basis argument establishing that the governing dimensional bound applies not to one representative operator but to the exhibited six-member local dimension-four parity-odd operator basis admitted by the minimal-coupling field content under stated construction rules (completeness of the basis is asserted from those rules, not proved by exhaustive symbolic enumeration).

Relation to the companion paper.—This survey builds on, and is a companion to, the paper establishing the algebraic torsion-elimination step and the resulting minimal axial–axial four-fermion contact interaction

* houston@hubify.com

(Route 1) in the same framework [1]. The division of content between the two papers is as follows. The companion publishes the Cartan elimination and the minimal contact operator, a regulated mean-field NJL check of the Route-1 condensate channel, the zero-spin-branch perturbation-transparency theorem, and a Fierz-rearrangement convention appendix with its released verification script; we cite those results rather than re-deriving them. The present survey contributes the material the companion does not publish: the systematic barrier catalog (Sec. III, Table I, Fig. 1), the Route-2/Route-3 amplitude-budget closures (Sec. IV), and the operator-basis argument (Sec. V) with its supporting appendices. On the Route-2/Route-3 mappings the two papers take one consistent position: the companion does not retain the historical one-loop Holst/Immirzi-running mappings to late-time dark energy as results, because the underlying calculations do not derive the required cosmological stress tensor and observable matching [2, 3]; this survey does not reinstate them — Sec. IV treats them as historical candidate routes and closes them as bounded amplitude budgets, without supplying (or needing) the missing stress-tensor derivation.

Contributions.—So that the publishable unit of this survey can be judged in one place: beyond the single-scale NDA argument of Appendix A and the companion’s published results, this paper contributes (i) the systematic route↔barrier taxonomy (Table I, Fig. 1) with its explicit three-tier evidentiary classification (Table II); (ii) the operator-basis closure of Sec. V / Appendix A 1 — the six-member genuine dimension-4 parity-odd basis {O1–O6} with its script-verified reduction identities (completeness of the basis is asserted from the stated construction rule, Sec. V, not proved); (iii) the integrated Route-3 bound $|\Delta\gamma/\gamma| \approx 1.4 \times 10^{-6}$ obtained from the Benedetti–Speziale one-loop flow; and (iv) the one-loop-grounded Route-2 amplitude budget. A self-contained statement and proof of the perturbation-transparency theorem (B14), the catalog’s sole Tier-I leg, is carried in Appendix D so that it can be refereed from this manuscript.

Sec. II fixes the minimal notation this survey needs. Sec. III presents the barrier catalog: the classification of results by evidentiary strength, the barrier-to-route structure (Fig. 1), the formal table (Table I), and the fourteen catalog entries. Sec. IV derives the Route-2 and Route-3 amplitude closures and ties all four enumerated routes together in a single evidentiary-status table. Sec. V presents the operator-basis argument. Sec. VI states explicitly what this survey does not claim. Sec. VII concludes.

II. CONVENTIONS AND SETUP

We work in minimal Einstein–Cartan–Holst (ECH) gravity: general relativity extended by promoting the spin connection to an independent field, with the Holst

term $\gamma^{-1} e^a \wedge e^b \wedge R_{ab}$ added to the first-order action with Barbero–Immirzi parameter γ . Signs, signature, and index conventions follow the companion paper’s setup [1]. We use natural units and $\kappa \equiv 8\pi G$ exactly; in terms of the *reduced* Planck mass $\overline{M}_{\text{Pl}} \equiv (8\pi G)^{-1/2} \simeq 2.44 \times 10^{18} \text{ GeV}$ this is $\kappa = \overline{M}_{\text{Pl}}^{-2}$, while order-of-magnitude prose throughout writes $\kappa \sim M_{\text{Pl}}^{-2}$ with the *full* Planck mass $M_{\text{Pl}} \equiv G^{-1/2}$ — a deliberate factor- 8π abuse (exactly, $\kappa = 8\pi M_{\text{Pl}}^{-2}$), never a definition.

Three convention flags are fixed here once. First, the imported one-loop results of Shapiró–Teixeira and Benedetti–Speziale (Sec. IV) are quoted in those papers’ own gravitational-coupling symbol, written $\tilde{\kappa}$ below, with $\tilde{\kappa}^2 \equiv 16\pi G$ — in their notation a single coupling of mass dimension -2 , *not* the square of this paper’s κ (under this paper’s κ , $\kappa^2 = M_{\text{Pl}}^{-4}$ carries dimension -4). Every $\tilde{\kappa}$ below refers to that imported convention. Second, the order-of-magnitude numerics in this survey use the full Planck mass $M_{\text{Pl}} \simeq 1.22 \times 10^{19} \text{ GeV}$; the reduced-versus-full distinction (2.44×10^{18} vs $1.22 \times 10^{19} \text{ GeV}$) is immaterial at the ≥ 58 -order margins quoted. Where a specific imported numeric depends on the choice, the usage is stated rather than hidden: the Table II R1 benchmark $\kappa n_\psi^2 / \rho_\Lambda \simeq 3.6 \times 10^{-69}$ is the companion’s published value [1], computed with the exact $\kappa = 8\pi G$ (equivalently the reduced-mass form) *and* the companion’s dark-energy normalization $\rho_\Lambda \approx (2.3 \text{ meV})^4 \approx 2.8 \times 10^{-11} \text{ eV}^4$; under the $\rho_\Lambda = (2.25 \text{ meV})^4$ input used in the Appendix A hierarchy, the same κn_ψ^2 evaluates to 3.9×10^{-69} — identical at the order-of-magnitude precision quoted (≈ 68 orders below ρ_Λ either way), while the Appendix A hierarchy $M_{\text{Pl}}^4 / \rho_{\Lambda, \text{obs}}$ uses the full mass; each is quoted at order-of-magnitude precision only. Third, the symbol pair (α, M) appears in two roles: in the operator ansatz of Sec. V [Eq. (5)], M denotes the LQG area-gap scale $M_{\text{area-gap}} \sim M_{\text{Pl}} / \sqrt{\gamma}$, while every *numerical* occurrence of the combination α/M in this survey (the Route-2 budget, the R4 channel, and Appendix A) uses the single R4-fitted anchor value $\alpha/M \sim 10^{-21} \text{ GeV}^{-1}$ (Sec. IV C); only the fitted combination α/M — never α or M separately — enters any quantitative statement, so no conclusion depends on resolving the two roles. Similarly, $\beta(\gamma)$ (always written with its argument) is the Route-2 renormalization-group function, while bare β is the observed birefringence angle.

Torsion in ECH is non-propagating: varying the first-order action with respect to the connection gives an algebraic (non-dynamical) constraint fixing the torsion tensor T^{abc} in terms of the fermion spin current S^{abc} , $T^{abc} = \kappa S^{abc}$. For minimally (totally antisymmetrically) coupled Dirac matter, the spin current is dual to the axial current, $S^{abc} = \frac{1}{4} \epsilon^{abcd} J_d^5$ with $J_d^5 = \bar{\psi} \gamma^\mu \gamma^5 \psi$. Substituting the algebraic constraint back into the action eliminates torsion in favor of a genuine four-fermion axial-axial contact interaction with coefficient fixed entirely by κ and γ ; the companion paper derives this elimination in full and obtains the resulting minimal contact operator, $-(3\kappa/16)[\gamma^2/(1+\gamma^2)](J^5)^2$, together with its mean-field

vacuum-condensate analysis (Route 1 of the four-route enumeration below) [1]. We do not repeat that derivation here; this survey begins from its result and extends the same setup to the remaining routes and to the barrier catalog that constrains all of them.

We enumerate four minimal-ECH channels through which the bounce-scale torsion sector could in principle source late-time dark energy: R1 (the NJL-type four-fermion contact interaction just described), R2 (one-loop graviton-sector corrections to the Holst term), R3 (quantum running of the Immirzi parameter γ), and R4 (a direct parity-odd coupling between the electromagnetic field and a spectator axion-like field or the fermion axial current, which also imprints cosmic birefringence). R1 is addressed in the companion paper; R2 and R3 are closed in Sec. IV below; R4 closes by a naturalness rather than an amplitude argument and is summarized, for completeness of the four-route accounting, in Sec. IV C. Its full derivation is outside the scope of this survey.

III. THE BARRIER CATALOG

Before cataloguing the individual barriers we fix their collective status precisely, so that their joint force is not overstated. We describe the catalog as *13 distinct mechanism-class constraints* (14 historical entries, with B8 subsumed by B14). Here ‘distinct’ and ‘mechanism-class’ mean that no barrier is a logical *consequence* of another and each probes a separate physical failure mode (amplitude suppression, thermal washout, operator decoupling, naturalness deficit, and so on); they do *not* assert that the barriers rest on disjoint assumptions or that each is an independent rigorous no-go theorem. In particular, several barriers share the same phenomenological on-shell scaling ansatz (Appendix A); several—B5 (scale separation), B6 (attractor sensitivity), B7 (parameter immunity), B10 (UV→IR specificity), and B13 (gravitational democracy)—are general naturalness or classification arguments that apply to broad classes of bounce/modified-gravity models rather than sharp ECH-specific calculations; and at least one, B9 (Liouville conservation), is an explicitly *heuristic* closure conditional on stated equilibrium assumptions (no particle production, no entropy injection) that realistic quantum bounces can violate. The catalog is therefore best read as a structured map of failure modes of mixed individual strength—quantitative amplitude bounds, naturalness arguments, and qualitative observations—whose collective value is the systematic coverage of the route space, not a claim that thirteen separately decisive theorems each independently exclude the framework. The two sharp, first-principles derivations behind the catalog are the Route-1 torsion-elimination derivation (a companion-paper result [1]) and the perturbation-transparency theorem (B14; established in the companion paper’s zero-spin-branch analysis and carried self-contained in Appendix D). Of these, only the perturbation-transparency

theorem enters the closure table (Table II) as a Tier-I leg: the torsion-elimination derivation is itself first-principles, but the Route-1 *closure* it feeds rests on Tier-II+III arguments, which is why the abstract credits exactly one Tier-I rigorous theorem. The remaining entries are constraints of the weaker classes just described and are labeled as such below.

We tested 7 foundation mechanism classes (Foundations A–G) and 6 additional observational channels (Branches H, J, L, M, N, O; the ECH perturbation gates enter through Branch H, whose first-principles closure B14 subsumes B8) for the possibility of connecting the ECH bounce to late-time dark energy. Each test yielded a named structural constraint. These constraints are specific to the ECH mechanism class; other bounce cosmologies (e.g., quintom scenarios) are not subject to them.

Constraint classification.—**Novel results** (Barriers 1, 2, 3, 4, 8, 10, 11, 12, 14): constraints first formulated within the ECH analysis, not immediate consequences of prior literature. Novelty here is a *provenance* label, not a rigor or specificity claim: eight of the nine are ECH-specific calculations, while B10 is a general naturalness dilemma first articulated in this catalog in the ECH-bridge context — consistent with its classification among the general naturalness arguments in the preamble above, in its own entry tag, and in Sec. VI. **Known results** (Barriers 5, 6, 7, 9): scale separation, attractor-sensitivity dilemma, parameter immunity, Liouville conservation (B9; heuristic ordering argument)—included to close mechanism classes that arise naturally in the ECH analysis. **Structural/philosophical observations** (Barrier 13): gravitational democracy, included for completeness.

A. Barrier details

The fourteen catalog entries of Table I are stated compactly below; the route each barrier constrains is given in brackets (the barrier→route map is drawn in Fig. 1). The entries are compact mechanism statements; several (B1, B4, B12) include explicit order-of-magnitude scaling relations inline.

B1 — Mass-Coupling Lock (Found. A) [R1]. In Poincaré gauge theory (PGT), ultralight torsion modes ($m_T \sim H_0$) require $g_{\text{eff}} \sim 1/(M_{\text{Pl}}\sqrt{|t_3|}) \sim H_0/M_{\text{Pl}} \sim 10^{-61}$, where t_3 is the quadratic-torsion coupling controlling the tensor-torsion mode mass, with $\sqrt{|t_3|} \sim m_T^{-1}$ for an ultralight mode (a labeled scaling ansatz of the PGT mass spectrum, not a derived equality). Achieving $g_{\text{eff}} \sim 1$ needs $m_T \sim M_{\text{Pl}}$; the required tuning $\delta m_T^2/m_T^2 \sim (H_0/M_{\text{Pl}})^2 \sim 10^{-122}$ is the standard cosmological-constant hierarchy.

B2 — Topological-Shift Duality (Found. B) [R1]. In metric-affine gravity, mass protection $\iff$ no geometric fingerprint: configurations protecting the pseudoscalar mass through topological structure eliminate the




FIG. 1. **Structure of the 14-entry barrier closure.** Top row: the 7 foundation mechanism classes (Foundations A–G; Barriers 1–7, mechanism-class barriers). Middle row: the 6 observational-channel branches (Branches H, J, L, M, N, O; Barriers 8–14, observational-channel barriers), grouped into four boxes — Branch H (B8, B14), Branch J (B9), Branches L/M (B10–B12), and Branches N/O (B13). Bottom row: the four closed dark-energy routes R1–R4; all four routes are closed, giving 13 distinct mechanism-class constraints. Arrows indicate which barrier classes constrain each route; Branch H carries arrows to all four routes because B14 (perturbation transparency) constrains R1–R4, while its co-resident B8 constrains R1 only — the in-figure labels record this per-barrier attribution (the upper Branch-H arrow to R1 carries both B8 and B14; the lower three-route fan is B14 alone), which is otherwise not resolved at box granularity. The branch letters are inherited unchanged from the historical catalog, which never assigned the letters I or K; no branch entries have been removed. B14 (the ECH perturbation-gate closure) sits in Branch H and subsumes B8 (parity-even interaction), so B8 is not counted separately; several entries share the scaling ansatz but each probes a distinct physical failure mode.

TABLE I. The 14 catalogue entries (13 distinct mechanism-class constraints; B8 subsumed by B14) on minimal ECH dark-energy routes. Note: Barriers 8 (parity-even interaction) and 14 (perturbation transparency) close the same observable channel (primordial-GW chirality / tensor parity violation) via related routes sharing the perturbation-transparency result; B14 is the first-principles theorem that subsumes B8 as the corresponding observational consequence. They are listed separately to preserve the historical mechanism-class catalog, but should not be counted as a separate mechanism-class constraint.

#	Barrier	Source	Mechanism Blocked
1	Mass-Coupling Lock	Found. A	Propagating torsion as DE
2	Topological-Shift Duality	Found. B	Geometric pseudoscalar mass protection
3	Scalar-Tensor Universality	Found. C	Distinctive geometric content on FRW
4	Planck Suppression	Found. D	Disformal / connection coupling effects
5	Scale Separation	Found. E	Global vacuum integral coupling
6	Attractor-Sensitivity Dilemma	Found. F	Initial-condition transfer to DE
7	Parameter Immunity	Found. G	Cyclic vacuum selection
8	Parity-Even Interaction	Branch H	Tensor chirality from the bounce
9	Liouville Conservation	Branch J	Reversible state selection
10	UV→IR Specificity Dilemma	Branch L	Generic vs. bounce-specific bridge
11	Decoupling Universality	Branch L/M	Light gauge field coupling
12	Vacuum Amplification Ceiling	Branch M	Gravitational wave amplitude
13	Gravitational Democracy	Branch N/O	Relics, baryogenesis, vacuum transitions
14	Perturbation Transparency	Branch H (ECH gates)	ECH-specific perturbation signatures

geometric content (the field reduces to a standard ALP after torsion elimination), while configurations preserving geometric content cannot protect the mass.

B3 — Scalar-Tensor Universality (Found. C) [R2]. On an FRW background, diffeomorphism invariance forces the torsion fluctuation to couple to the curvature invariants like any other scalar, with no ECH-specific observable; torsion decouples from the FRW background precisely at the bounce density, yielding no distinctive perturbation signal.

B4 — Planck Suppression (Found. D) [R2]. Disformal torsion couplings are Planck-suppressed by m_ϕ^2/M_{Pl}^2 or $(\partial\phi)^2/M_{\text{Pl}}^4$; at cosmological scales ($m_\phi \sim H_0$) these are $\mathcal{O}(10^{-122})$ —observationally inaccessible.

B5 — Scale Separation (Found. E) [R3]. The global vacuum integral $\int d^4x \sqrt{-g} \rho_\Lambda$ cannot be connected to the local bounce density without a mechanism to store and transfer the integrated vacuum energy across $N_{\text{tot}} \approx 92\text{--}94$ e -folds of inflation (Appendix A); none exists within minimal ECH. (General naturalness/classification

argument, not an ECH-specific calculation.)

B6 — Attractor-Sensitivity Dilemma (Found. F) [R3]. If post-bounce inflation converges to an attractor, bounce initial conditions are washed out; if it is sensitive to them, inflation destabilizes. The bounce cannot simultaneously seed dark energy *and* preserve standard inflation dynamics. (General naturalness argument.)

B7 — Parameter Immunity (Found. G) [R4]. Cyclic vacuum selection requires γ to vary across cycles, but γ is fixed by the LQG area spectrum at a universal value; LQG provides no landscape of γ values for selection to operate on. (General classification argument.)

B8 — Parity-Even Interaction (Branch H). [R1; subsumed by B14.] The spin-torsion interaction $(J^5)^2$ is parity-even (a product of two axial currents is a Lorentz scalar, not a pseudoscalar), so it cannot generate tensor chirality in primordial gravitational waves; independently confirmed by the perturbation-transparency result (B14).

B9 — Liouville Conservation (Branch J) [R2]. Phase-space volume conservation prevents irreversible selection among post-bounce states, closing the “vacuum selection at the bounce” class: the time-symmetric bounce selects no net dark-energy state. This is a *heuristic* closure under explicit assumptions (closed non-dissipative evolution, no particle production, no entropy injection); dissipative/particle-producing bounces evade B9 and are constrained instead by B10–B11, so B9 is never used as a stand-alone closure.

B10 — UV→IR Specificity Dilemma (Branch L) [R3]. Any bridge from Planck-scale bounce physics to the late-time H_0 scale must be either generic (explaining *any* vacuum energy) or bounce-specific (requiring free parameters equivalent to the cosmological constant); no ECH mechanism achieves both. (General naturalness argument.)

B11 — Decoupling Universality (Branches L/M) [R3]. At low energies all gauge fields decouple from the Planck-suppressed torsion sector equally; the ECH bounce cannot preferentially couple to photons or dark-energy degrees of freedom without new non-minimal couplings.

B12 — Vacuum Amplification Ceiling (Branch M) [R3]. Gravitational-wave production from the ECH bounce is bounded above by $\Omega_{\text{GW}}^{\text{ECH}}|_{\text{bounce}} \lesssim (\rho_{\text{crit}}/\rho_{\text{Pl}})^2 \simeq 0.07\text{--}0.17$, using the LQG-bounce critical-density window $\rho_{\text{crit}}/\rho_{\text{Pl}} \simeq 0.27\text{--}0.41$: the 0.41 endpoint is the canonical effective-LQC value quoted by Ashtekar–Singh at the standard area-gap choice $\gamma = 0.2375$ [4], while the 0.27 endpoint substitutes the SU(2) black-hole-entropy value $\gamma \approx 0.274$ into the same $\rho_{\text{crit}} = \sqrt{3}/(32\pi^2\gamma^3)\rho_{\text{Pl}}$ formula — an internal extrapolation across entropy-counting schemes established in the companion paper [1], not a value quoted in Ref. [4] — so

the window is scheme-dependent rather than a published LQC range. The quadratic scaling is an order-of-magnitude ceiling *ansatz* (not derived here); B12 is used only as a global ceiling. This bounce-epoch fraction is not directly comparable to the present-day PTA spectral density $\Omega_{\text{GW}}(f_{\text{nHz}}) \sim 10^{-9}$ (which differs by redshift/dilution and by frequency-integration); a quantitative NANOGrav comparison requires propagating the bounce GW spectrum through the transfer function to the nHz band, deferred to a dedicated bounce-GW paper, so B12 closes as a global energy-density-fraction ceiling, not a direct NANOGrav exclusion.

B13 — Gravitational Democracy (Branches N/O) [R4]. Torsion couples democratically to all spin-1/2 species and cannot preferentially source baryogenesis, dark-matter relics, or vacuum transitions without species-dependent non-minimal couplings. (Structural/philosophical observation, included for completeness.)

B14 — Perturbation Transparency (Branch H, ECH perturbation gates). [R1–R4.] For canonical scalar matter, torsion vanishes at all perturbation orders and the Holst sector decouples from all scalar/tensor perturbation observables. The statement and proof are carried self-contained in Appendix D; the tensor-sector extension and the explicit second-order Holst-term verification are elaborated in the companion paper’s zero-spin-branch analysis [1]. This is the catalog’s sole Tier-I closure leg (Table II).

IV. THE ROUTE-2/ROUTE-3 DARK-ENERGY ROUTE CLOSURES

The two routes closed in this section are *historical* candidate mappings — one-loop Holst-sector corrections (Route 2) and Immirzi-parameter running (Route 3) mapped onto a late-time dark-energy density — that the companion paper explicitly does not retain as results, for the reason recorded in Sec. I. Nothing in this section reinstates them. The closures below are upper-bound *amplitude budgets* demonstrating that any completion of these mappings falls short of the observed signal in its channel by tens of orders of magnitude — for Route 2 the budget is evaluated against the *observed birefringence amplitude* [Eq. (2)], for Route 3 against the *observed dark-energy density* — which is precisely why the catalog closes the corresponding channels. The published one-loop inputs enter as normalization anchors for the budgets, not as derivations of a cosmological observable.

A. Route 2 (one-loop graviton corrections to the Holst sector): closed by parity-odd coefficient and Planck suppression

At the classical level the Holst term $\gamma^{-1} e^a \wedge e^b \wedge R_{ab}$ vanishes identically in the torsion-free sector by the first

algebraic Bianchi identity (it is *not* a topological invariant in the sense of the Pontryagin density; it reduces to the Nieh–Yan density plus a torsion-squared piece, both of which vanish at $T = 0$, as established in the companion paper’s zero-spin-branch analysis [1]) and reduces to the Nieh–Yan density on shell once torsion is integrated out [5, 6]. Quantum corrections from minimally coupled fermions, however, generate a parity-odd coupling between the gravitational field and the chiral fermion current at one loop, with the Holst sector developing running couplings analyzed via renormalization-group methods in Einstein–Cartan + Holst gravity [2, 7]. Motivated by (but *not literally derived in*) the Holst+non-minimal-fermion construction of Mercuri and the one-loop analysis of Shapiro & Teixeira—those works establish the classical structure of the Holst term coupled to fermions, the Nieh–Yan invariant, and the one-loop running of the Holst sector, not this exact effective operator—we adopt the phenomenological one-loop parity-odd operator

$$\Gamma_{\text{one-loop}}^{\text{parity-odd}} = -\frac{1}{16\pi^2} \frac{\beta(\gamma)}{M_{\text{Pl}}} \int d^4x \sqrt{-g} \partial_\mu \vartheta_{\text{NY}}(x) J^{5\mu}, \quad (1)$$

where $\vartheta_{\text{NY}}(x)$ is the Nieh–Yan pseudoscalar (mass dimension +1; written ϑ_{NY} to distinguish it from a photon-sector spectator ALP θ of the R4 channel, Sec. IV C; no identification between the two fields is assumed anywhere in this paper), $J^{5\mu}$ is the fermion axial current, and $\beta(\gamma)$ is a slowly varying function of γ . The parity classification of Eq. (1) — the operator is intrinsically parity-even; the parity violation is background-induced — and the anomaly-mediated photon chain used in the dimensionless ratio below are stated in Appendix B. The function $\beta(\gamma)$ is written with its explicit argument throughout to distinguish this renormalization-group function from the cosmic birefringence angle β of the R4 channel (Sec. IV C). The coefficient structure of Eq. (1) is *grounded in the explicit one-loop computation* of Shapiro & Teixeira [2], who renormalize on-shell Einstein–Cartan+Holst gravity with external vector and axial fermion currents. The parity-odd, Immirzi-dependent current coupling they renormalize is $\lambda_4 = \gamma \tilde{\kappa}^2 (W \cdot J)$ (their Eq. 37; $\tilde{\kappa}^2 \equiv 16\pi G \sim M_{\text{Pl}}^{-2}$, the imported convention fixed in Sec. II), with classical coefficient $\alpha_4 = -6/(1 + \gamma^2)$ (their Eq. 41) and one-loop divergence coefficients $\Omega_{44} = -(378 + 783\gamma^2)/[20(1 + \gamma^2)^2]$, $\Omega_{24} = -81\gamma^2/(1 + \gamma^2)^2$ (their Eq. 42, arXiv version; the λ_4 flow, their Eq. 51, is driven by the $\Omega_{44}, \Omega_{24}, \Omega_{34}$ family). Their master renormalization-group equation $d\lambda/dt = -\sigma/(4\pi)^2$ (their Eq. 46) carries *exactly* the $1/(16\pi^2)$ loop factor already written in Eq. (1). Three features of the Route-2 budget are therefore fixed by a published one-loop result rather than adopted: (i) the loop factor $1/(16\pi^2)$; (ii) the coefficient $\beta(\gamma)$ is an $\mathcal{O}(1)$ rational function of the Immirzi parameter [the α_4/Ω_{4x} family, e.g. $|\Omega_{44}/\alpha_4| = (378 + 783\gamma^2)/[120(1 + \gamma^2)]$, which is ≈ 3.3 at the LQG value $\gamma \approx 0.24$ and $\mathcal{O}(1)$ – $\mathcal{O}(5)$ across $\gamma \lesssim \mathcal{O}(1)$], not a free normalization; and (iii) the explicit $\tilde{\kappa}^2 \sim M_{\text{Pl}}^{-2}$ on every renormalized charge, i.e. the

Planck suppression in the prefactor $\beta(\gamma)/M_{\text{Pl}}$. What the one-loop analysis does *not* fix is the single *absolute* normalization: Shapiro & Teixeira show the coupled flow for $\lambda_4(t)$ and $\gamma(t)$ (their Eqs. 51, 58) is a Riccati system whose particular-solution roots are complex for real γ , so the Shapiro–Teixeira λ_4 – γ system *has no renormalization-group fixed point* (in contrast to the distinct Benedetti–Speziale flow of Route 3 below, which has one) and they were “unable to solve it in a completely satisfactory way.” A fully-derived Route-2 amplitude would require a UV boundary condition plus a controlled solution of that non-perturbative flow (or a dedicated matching calculation for the exact $\partial_\mu \vartheta_{\text{NY}} J^{5\mu}$ operator). We therefore treat the *absolute* $\beta(\gamma)$ normalization as a bounded EFT input while the loop factor, Immirzi-rational coefficient, and Planck suppression are one-loop-grounded. This is strictly weaker than the Route-3 result below (whose Benedetti–Speziale β -function integrates cleanly to $|\Delta\gamma/\gamma| \approx 1.4 \times 10^{-6}$) but strictly stronger than a free ansatz: the residual freedom is a single $\mathcal{O}(1)$ normalization, and the off-shell dimension- $(+1)$ parity-odd operator is in any case bounded by the single-scale NDA no-go of App. A regardless of that $\mathcal{O}(1)$ coefficient, since ST’s explicit $\tilde{\kappa}^2 \sim M_{\text{Pl}}^{-2}$ confirms the missing powers are Planck powers, not a light scale. Because the closure below retains ≈ 60 (conservatively ≥ 58) orders of margin, it is insensitive to this residual $\mathcal{O}(1)$ ambiguity in $\beta(\gamma)$. The dimensionless coefficient is $\mathcal{O}(\alpha_{\text{em}}/4\pi)$ multiplied by the Planck mass to a single negative power. Explicitly, the prefactor $\beta(\gamma)/M_{\text{Pl}}$ carries mass dimension -1 (the division by M_{Pl} , not multiplication): with ϑ_{NY} of dimension $+1$, $\partial_\mu \vartheta_{\text{NY}}$ has dimension $+2$ and the axial current $J^{5\mu} = \bar{\psi} \gamma^\mu \gamma^5 \psi$ has dimension $+3$, so the full term — the dimension- $(+5)$ integrand $\partial_\mu \vartheta_{\text{NY}} J^{5\mu}$ times the dimension- (-1) prefactor $\beta(\gamma)/M_{\text{Pl}}$ written outside the integral in Eq. (1) — carries dimension $-1 + 2 + 3 = +4$ and the action is dimensionless, as required. Once the dimension- $(+2)$ background value $\partial_\mu \vartheta_{\text{NY}} \sim H_0^2 \sim 10^{-66} \text{ eV}^2$ at the present epoch is substituted (the cosmological Nieh–Yan pseudoscalar, of dimension $+1$ and Hubble-scale amplitude, evolving on the Hubble time), the induced birefringence accumulated between recombination and today is, in the dimensionless ratio $\Delta\theta_{\text{one-loop}}/\Delta\theta_{\text{obs}}$,

$$\begin{aligned} \frac{\Delta\theta_{\text{one-loop}}}{\Delta\theta_{\text{obs}}} &\sim \frac{\alpha_{\text{em}}}{4\pi} \frac{H_0/M_{\text{Pl}}}{M_{\text{Pl}}(\alpha/M)\beta_{\text{obs}}} \\ &\sim \frac{\alpha_{\text{em}}}{4\pi} \frac{H_0}{M_{\text{Pl}}} \frac{M}{\alpha M_{\text{Pl}}\beta_{\text{obs}}}, \end{aligned} \quad (2)$$

where $\alpha_{\text{em}}/(4\pi) \approx 6 \times 10^{-4}$ (more precisely 5.8×10^{-4}), rounded *up* to 10^{-3} in the explicit order-of-magnitude substitutions below — an overestimate of the numerator and hence conservative for the suppression claim; the closure is robust to this factor. Here $H_0/M_{\text{Pl}} \sim 10^{-61}$, and the R4-fitted coupling $\alpha/M \sim 10^{-21} \text{ GeV}^{-1}$ (Sec. IV C) gives $M_{\text{Pl}} \cdot (\alpha/M) \sim 10^{19} \text{ GeV} \cdot 10^{-21} \text{ GeV}^{-1} = 10^{-2}$. Plugging in $\beta_{\text{obs}} = 0.342^\circ \approx 6 \times 10^{-3} \text{ rad}$ ($0.342 \times \pi/180 = 5.97 \times 10^{-3}$), the dimensionless ratio is

$\Delta\theta_{\text{one-loop}}/\Delta\theta_{\text{obs}} \sim 10^{-3} \cdot 10^{-61}/(10^{-2} \cdot 6 \times 10^{-3}) \approx 10^{-60}$ (canonical evaluation of the displayed contraction; we conservatively allow up to two orders of magnitude for unmodeled higher-order loop-ordering corrections, i.e. suppression by at least 10^{-58} — an explicit conservatism allowance, *not* a derived range; the eV-vs-GeV unit conversion is exact $1 \text{ GeV} = 10^9 \text{ eV}$ and is not a source of ambiguity). The role of each denominator factor is stated explicitly so the contraction is not misread as a double division: the displayed ratio normalizes the one-loop amplitude *both* by the observed angle β_{obs} *and* by the dimensionless strength $M_{\text{Pl}}(\alpha/M) \sim 10^{-2}$ of the R4-fitted coupling that reproduces that angle — i.e. it states the suppression relative to the fitted-coupling benchmark, the weaker (more conservative) normalization. Contracting directly against the observed angle alone gives $(\alpha_{\text{em}}/4\pi)(H_0/M_{\text{Pl}})/\beta_{\text{obs}} \approx 10^{-3} \cdot 10^{-61}/(6 \times 10^{-3}) \approx 2 \times 10^{-62}$, *two additional orders* of suppression; the quoted $\sim 10^{-60}$ is therefore the conservative side of this bookkeeping choice, and the closure holds a fortiori under either contraction. Thus the one-loop induced β is suppressed by ≈ 60 (conservatively ≥ 58) orders of magnitude relative to the observed signal. This is the survey's one authoritative Route-2 robustness statement: even inflating the one-loop coefficient by ten orders of magnitude leaves $\gtrsim 48$ orders of suppression margin, so no plausible $\mathcal{O}(1)\text{--}\mathcal{O}(10^{10})$ normalization ambiguity reopens the route. We adopt this contraction as the canonical Route-2 estimate; an alternative ordering that contracts the H_0 factor with the dimensionful coupling differently yields a deliberately loose $\sim 10^{-33}$ upper bound, not used in the closure. The canonical-bound conclusion that the one-loop induced β is amplitude-suppressed many orders of magnitude below the observed WMAP+Planck birefringence signal (with ACT DR6 follow-up) is robust to this choice; Route 2 remains exploratory framing, not load-bearing for the no-go. The Route-2 amplitude is therefore far below not only the WMAP+Planck birefringence sensitivity but the observed central value itself; the one-loop Holst-sector parity-odd term cannot account for the observed birefringence amplitude. (A naive comparison of a rotation rate $\dot{\beta}$ in eV against an angle uncertainty in eV would silently treat eV's as dimensionless; the dimensionless reduction above avoids this and recovers the standard R2 amplitude-suppression closure.) *Closure: amplitude-suppressed by M_{Pl}^{-1} and one-loop factor $\alpha_{\text{em}}/(4\pi)$.* The evidentiary status of the estimate, stated once: Eq. (1) is an *illustrative upper-bound amplitude budget* constructed as the natural EFT operator at the $M_{\text{Pl}}^{-1}\alpha_{\text{em}}/(4\pi)$ scale, *not* a result extracted from Mercuri [7] or Date–Kaul–Sengupta [8], which establish the classical Holst/Nieh–Yan structure but not this coefficient; it is presented as an amplitude-budget bound, not a rigorous derivation, and the same logic applies to Route 3 below.

B. Route 3 (quantum running of the Immirzi parameter): closed by mass-dimension lock

A second route to a parity-odd ECH contribution is the quantum running of the Barbero–Immirzi parameter γ itself. Date, Kaul & Sengupta analyzed the Holst term coupled to fermions and the Nieh–Yan invariant in the chiral-matter setting [8]; that analysis establishes a topological interpretation of γ and motivates a γ -running in the presence of chiral asymmetry, but does *not* itself present the explicit RG equation used below. Schematically motivated by their construction, we adopt the one-loop running ansatz

$$\frac{d\gamma}{d\ln\mu} = \frac{1}{12\pi^2} (N_F^L - N_F^R) \gamma + \mathcal{O}(\gamma^2), \quad (3)$$

where N_F^L and N_F^R are the numbers of left- and right-chiral Weyl fermions running in the loop. The $1/(12\pi^2)$ prefactor is the natural chiral-loop coefficient at this order; we use Eq. (3) only as an upper-bound EFT ansatz for the Route-3 amplitude budget and do not claim it is taken verbatim from [8]. The actual fermion-induced perturbative running of the Immirzi parameter is computed by Benedetti & Speziale [9], who find a β -function whose sign depends on $|\gamma|$ through four-fermion interactions generated when fermions are coupled to the Holst sector; our Eq. (3) is a chiral-count EFT bound rather than the full perturbative result, and is used solely for the amplitude budget below. We can, moreover, do better than an estimate. The actual fermion-coupled one-loop β -function was computed by Benedetti & Speziale [3] (their Eq. 7),

$$\mu \frac{\partial \gamma^2}{\partial \mu} = -(\gamma^2 - 1) \frac{\mu^2 \tilde{\kappa}^2}{(8\pi)^2} (23\gamma^2 + 5), \quad \tilde{\kappa}^2 \equiv 16\pi G, \quad (4)$$

whose only fixed point (of this Benedetti–Speziale flow; the fixed-point-free Riccati system above is Shapiro–Teixeira's Route-2 $\lambda_4\text{--}\gamma$ system, a different flow) is the ultraviolet-attractive $\gamma^2 = 1$ (formally outside perturbative control), with the sign of the flow set by whether $|\gamma| \gtrless 1$ and the running driven by the radiatively-generated four-fermion interaction. The decisive feature is the explicit $\mu^2 \tilde{\kappa}^2 \simeq (\mu/M_{\text{Pl}})^2$ prefactor: the running is *power-suppressed* by the renormalization scale in Planck units, not logarithmic, so the accumulated $\int \beta_{\gamma^2} d\ln\mu \propto \int \mu d\mu$ is dominated by the ultraviolet endpoint and is of order $(\mu_{\text{UV}}/M_{\text{Pl}})^2$. Integrating Eq. (4) numerically from a GUT-scale ultraviolet boundary $\mu_{\text{UV}} \sim 10^{16} \text{ GeV}$ down to $\mu_{\text{IR}} \sim 1 \text{ GeV}$ (with γ near the LQG value $\gamma \approx 0.24$) gives, in agreement with a frozen-coefficient analytic estimate, $|\Delta\gamma/\gamma| \approx 1.4 \times 10^{-6}$. This is a genuine integrated running, *not* an ansatz bound: the explicit $\mu^2 \tilde{\kappa}^2 \simeq (\mu/M_{\text{Pl}})^2$ prefactor makes the flow power-suppressed, so $|\Delta\gamma/\gamma| \sim (\mu_{\text{UV}}/M_{\text{Pl}})^2$ and the integral is dominated by the ultraviolet endpoint. The result is perturbatively controlled for any sub-Planckian UV boundary; only a literal Planck-scale boundary would push $|\Delta\gamma/\gamma| \rightarrow \mathcal{O}(1)$,

precisely the regime Benedetti & Speziale flag as outside perturbative control (the $\gamma^2 = 1$ fixed point corresponds to a divergent four-fermion coupling), so a Planck-scale evaluation would require UV-completion input beyond the one-loop β -function. The derived physical running therefore only *strengthens* Route 3, and this GUT-UV integrated value is the primary Route-3 estimate. Propagated to the dark-energy channel through the paper’s own $(\Delta\gamma/\gamma) \cdot (H_0/M_{\text{Pl}})$ mass-dimension suppression, the derived torsion/Immirzi contribution to ρ_Λ sits ~ 61 – 67 orders of magnitude below the observed dark-energy density $\rho_{\Lambda, \text{obs}} \approx (2.25 \text{ meV})^4$ — so the no-go closes with enormous margin, now from a *derived* integrated result rather than an ansatz upper bound. We nonetheless retain the far larger chiral-count estimate Eq. (3), $\Delta\gamma/\gamma \sim 0.3$, as a deliberately pessimistic *upper* bound in the budget below. In the Standard Model, the chiral asymmetry is generated by the $SU(2)_L$ doublets. Integrating Eq. (3) over the running gives $\Delta\gamma/\gamma \approx (N_F^L - N_F^R) \ln(\mu_{\text{GUT}}/\mu_{\text{IR}})/(12\pi^2)$; with a net chiral count $(N_F^L - N_F^R) = \mathcal{O}(1)$ and a GUT-to-IR lever arm $\ln(\mu_{\text{GUT}}/\mu_{\text{IR}}) \approx 30$ – 37 ($\mu_{\text{GUT}} \sim 10^{16}$, $\mu_{\text{IR}} \sim 1 \text{ GeV}$, $\ln 10^{16} \approx 36.8$), this is numerically $\Delta\gamma/\gamma \approx 0.25$ – 0.31 (a few $\times 10^{-1}$; e.g. $32/(12\pi^2) \approx 0.27$). We therefore adopt the larger $\Delta\gamma/\gamma \sim 0.3$ as the *conservative* (least-suppressed) order-of-magnitude estimate — *not* a precisely derived value, but a conservative upper bound consistent with the $|\gamma|$ -dependent Benedetti–Speziale β -function structure recorded above (four-fermion-driven, sole fixed point at $\gamma^2 = 1$, sign set by $|\gamma| \gtrless 1$) — and note that the Route-3 closure below is insensitive to its precise value: even this $\mathcal{O}(0.3)$ running is $\ll 1$, so γ retains its order of magnitude, and it still leaves $\gtrsim 60$ orders of suppression margin. The Holst sector amplitude that this running can source is fixed by mass dimension: any operator built from γ , R_{ab} , e^a , and the chiral current $J^{5\mu}$ must carry dimension four, which forces a single power of M_{Pl}^{-1} in the prefactor in any cosmologically relevant scalar-curvature regime. Plugging the conservative $\Delta\gamma/\gamma \sim 0.3$ into the resulting parity-odd amplitude, the cosmologically integrated effect is suppressed by an additional factor of $(\Delta\gamma/\gamma) \cdot (H/M_{\text{Pl}}) \sim 3 \times 10^{-62}$ relative to the dimensionless parity-odd amplitude budget associated with a dark-energy-scale source, closing this route by many orders of magnitude. *Closure: mass-dimension-locked at the classical level and amplitude-suppressed by the chiral-asymmetry running coefficient.* *Ansatz vs derivation (R2/R3):* R3’s magnitude is now a *derived integrated result*: for a sub-Planckian (GUT-scale) UV boundary the verified Benedetti–Speziale one-loop β -function Eq. (4) integrates to $|\Delta\gamma/\gamma| \approx 1.4 \times 10^{-6}$, replacing the prior chiral-count ansatz upper bound as the primary Route-3 estimate. R2’s amplitude coefficient is now *one-loop-grounded*: its loop factor $1/(16\pi^2)$, $\mathcal{O}(1)$ Immirzi-rational coefficient, and $\tilde{\kappa}^2 \sim M_{\text{Pl}}^{-2}$ Planck suppression are fixed by the Shapiro & Teixeira [2] one-loop renormalization of the Holst+fermion sector (their Eqs. 41–42, 46); only the single absolute normalization

remains a bounded EFT input, because the Shapiro–Teixeira λ_4 – γ flow has no fixed point and is not perturbatively solvable in closed form, and the operator is bounded by the single-scale NDA no-go regardless of that $\mathcal{O}(1)$ normalization. The kept ansatz value Eq. (3) ($\Delta\gamma/\gamma \sim 0.3$) is retained only as a deliberately pessimistic upper bound; even under it the amplitude suppression is $\sim 3 \times 10^{-62}$ vs ρ_Λ , and the derived 1.4×10^{-6} value only widens that margin. *Strict theoretical limitation:* the NDA one-loop operator (R2) and Immirzi-running (R3) inputs used here are bounded EFT inputs, not UV-complete flow derivations—the R2 absolute normalization requires a controlled solution to the ST Riccati system (no real fixed point; ST explicitly state they could not solve it satisfactorily), and the R3 chiral-count Eq. (3) is an order-of-magnitude ansatz rather than the full Benedetti–Speziale four-fermion-driven RG flow at all loop orders. These represent the strict theoretical limitation of the R2/R3 analysis within current loop-gravity technology; the ≈ 58 – 60 -order (and larger) closure margins make the qualitative conclusions robust to this limitation, but the inputs cannot be promoted to precision derivations without a UV-complete treatment of the coupled Holst+fermion RG system.

C. The fourth channel and the combined closure

The fourth enumerated channel, R4, is a direct parity-odd coupling between the electromagnetic field and a spectator axion-like field or the fermion axial current, $\mathcal{L}_{\text{CS}} \supset -\frac{1}{4}(\alpha/M) \phi \tilde{F}_{\mu\nu} F^{\mu\nu}$, producing a rotation of the CMB polarization plane (cosmic birefringence). Fitting this coupling to the observed WMAP+Planck birefringence signal $\beta_{\text{obs}} = 0.342^\circ \pm 0.094^\circ$ [10, 11] (comparable to the independent ACT DR6 measurement $\beta = 0.215^\circ \pm 0.074^\circ$ [12]) gives the anchor value $\alpha/M \sim 10^{-21} \text{ GeV}^{-1}$ used above as a numerical input to the Route-2 amplitude estimate. The algebraic origin of that scale, carried from the companion’s derivation [1], is the standard small-rotation result $\beta = (\alpha/2M) \Delta\phi_{\text{rec} \rightarrow \text{today}}$ (the rotation angle is half the coupling times the coherent field excursion), with excursion $\Delta\phi \sim \sqrt{2\rho_\theta}/m_\theta \sim M_{\text{Pl}}$ for a frozen ($m_\theta \lesssim H_0$) spectator carrying $\rho_\theta \sim \rho_\Lambda$; matching $\beta_{\text{obs}} \approx 6 \times 10^{-3} \text{ rad}$ then gives $\alpha/M \sim 2\beta_{\text{obs}}/M_{\text{Pl}} \sim 10^{-21} \text{ GeV}^{-1}$. R4 does not close by amplitude suppression: the same coupling that reproduces β_{obs} can, with a free normalization, also be tuned to reproduce ρ_Λ , but minimal ECH does not derive either the required ultralight mass $m_\theta \sim H_0$ or the fitted α/M from first principles—so the channel closes as a naturalness/explanatory-deficit objection (it relocates rather than solves the cosmological-constant problem) rather than as an amplitude exclusion. Its full derivation, and the companion consistency check against a spectator-ALP birefringence benchmark, are outside the present survey’s scope [1, 13].

Within the four-channel enumeration of Sec. II, Routes

R1–R4 cover the four parity-odd/dark-energy channels of this framework. Two operators occasionally raised as omissions from this four-channel enumeration — the Jackiw–Pi gravitational Chern–Simons term $R \wedge \tilde{R}$ and the parity-odd four-fermion partner of R1 carrying the $\gamma_{\text{BI}}/(\gamma_{\text{BI}}^2+1) \cdot 8\pi G$ coefficient — are closed in the companion analysis [1] (total derivative for constant coupling and R4-class otherwise, for the Jackiw–Pi term; Planck suppression plus vanishing mean field, inheriting R1, for the four-fermion partner); the dimension-6 parity-odd four-fermion basis is Fierz-closed and uniformly Planck-suppressed by the projection lemma of Appendix C, so only a non-minimal completion lies outside the established no-go. Within the four enumerated channels: R1 (NJL contact) is amplitude-suppressed by M_{Pl}^{-2} and parity-even. R2 (one-loop graviton corrections) is amplitude-suppressed by M_{Pl}^{-1} and the one-loop factor $\alpha_{\text{em}}/(4\pi)$. R3 (Immirzi running) is mass-dimension-locked and additionally suppressed by the chiral-asymmetry beta function. R4 (parity-odd CMB coupling) is closed by a naturalness objection: with α/M treated as a free parameter, a spectator-ALP fit reproduces both β_{obs} and ρ_Λ , but minimal ECH does not derive $m_\theta \sim H_0$ or the fitted α/M , so the channel closes at the level of an explanatory deficit rather than an amplitude exclusion.

Tiered closure structure: R1–R3 achieve *mathematical amplitude closure* (the predicted amplitude is suppressed by explicit M_{Pl} powers and loop factors, placing it many orders below any observable threshold regardless of normalisation assumptions); R4 achieves *naturalness/explanatory closure* (the amplitude can be matched to observations but only at the cost of importing the cosmological-constant fine-tuning the channel was meant to resolve). These are logically distinct closure modes and are classified separately in the evidentiary table below.

a. Evidentiary status of each leg. So that the strength of the closure is not overread, we state explicitly the evidentiary level at which each route and each structural result is established. We use a three-tier scale: **(I) rigorous result** — a deductive consequence of stated equations/identities, holding exactly within an explicitly bounded scope; **(II) structural argument** — a qualitative or order-of-magnitude argument from established physics (symmetry, parity, naturalness) that does not turn on a fitted number; and **(III) ansatz-level dimensional estimate** — an amplitude budget evaluated under an explicitly-labeled scaling ansatz, honest to a factor but not a first-principles derivation. Table II classifies every leg on this scale. No leg is claimed at a level higher than this table records: in particular, the only Tier-I (rigorous) leg is the perturbation-transparency result for canonical scalar matter; R2 is a Tier-III ansatz-level estimate, and the R1/R3 *amplitude* legs are likewise Tier-III (their structural components — the parity-even mean-field fact for R1 and the mass-dimension lock for R3 — are the Tier-II entries Table II records alongside them);

and the R4 closure is a Tier-II *naturalness/explanatory-deficit* objection, *not* an amplitude exclusion and conditional on the on-shell scaling ansatz of Appendix A. The aggregate “four-route channel-level closure” is therefore exactly that: a channel-level (not operator-level) statement whose individual legs sit at the levels tabulated, reinforced but not promoted by the mechanism-class catalog of Sec. III.

The phenomenological parameter α/M introduced above is therefore best understood as the R4-bounded coupling required to match β_{obs} , with the dark-energy density supplied by an unrelated sector (e.g. a quintessence field [14], a quintom-class two-field scenario [15], or by a positive cosmological constant), rather than by ECH-internal physics. This inversion of the prior mass-coupling lock is a core structural finding reused throughout this catalog.

V. THE OPERATOR-BASIS ARGUMENT

The barrier catalog and the Route-2/Route-3 closures above bound the amplitude of *one* representative parity-odd operator. A referee may reasonably ask whether the same conclusion survives once the full local operator basis is considered. This section answers that question: we exhibit the finite basis of local, gauge-invariant, diffeomorphism-covariant parity-odd densities of mass dimension exactly four admitted by the minimal-ECH field content under the construction rule stated below, and show that the single-scale dimensional bound applies to every member of that basis, not to one representative.

As a phenomenological ansatz for a direct parity-odd curvature coupling — motivated by, but not literally derived from, the Holst-plus-non-minimal-fermion construction of Mercuri [7] — consider the parity-odd term

$$S_{\text{eff}} = \frac{\alpha}{M} \int e_I \wedge e_J \wedge \mathcal{F}^{IJ}[K, \hat{R}], \quad (5)$$

where $M = M_{\text{area-gap}} \sim M_{\text{Pl}}/\sqrt{\gamma}$ is the LQG area-gap mass scale (from the LQG area-gap $\Delta \propto \gamma \ell_P^2$, the inverse-length / mass scale is $M_\Delta \sim M_{\text{Pl}}/\sqrt{\gamma}$ up to numerical constants), α is a dimensionless coupling, and $\mathcal{F}^{IJ}[K, \hat{R}]$ denotes the curvature two-form of the full (torsionful) connection, written as a functional of the contorsion K and the Levi-Civita curvature $\hat{R}$; the component form Eq. (6) displays its leading contribution.

In components, the leading contribution reduces to

$$S_{\text{eff}} = \int d^4x \sqrt{-g} \frac{\alpha}{M} \epsilon^{\mu\nu\rho\sigma} e_\mu^I e_\nu^J \mathcal{F}_{IJ\rho\sigma}, \quad (6)$$

which has naive mass dimension $[\mathcal{L}_{\text{odd}}] = +1$ —three units short of the required $+4$ (with $[\alpha] = 0$, $[M] = +1$, $[e_\mu^I] = 0$, and $[\mathcal{F}_{IJ\rho\sigma}] = +2$, the integrand carries mass dimension $-1+2 = +1$). We state this off-shell mismatch deliberately, and we address it immediately by identifying the physical foundation before using the shorthand.

TABLE II. Evidentiary status of each closure leg, on the three-tier scale of the text: (I) rigorous result within stated scope, (II) structural argument, (III) ansatz-level dimensional estimate. The table records the highest level at which each leg is claimed; no leg is asserted more strongly elsewhere. “Rules out” is to be read at channel-amplitude granularity, not as an operator-level theorem.

Leg	What it rules out (channel-amplitude level)	Evidentiary status
Perturbation transparency (companion paper [1])	The Holst sector / Barbero–Immirzi parameter contributing to scalar or tensor perturbation observables, for canonical scalar matter	(I) Rigorous within stated scope: $T = 0$ follows from zero scalar spin density, and the Holst dual contraction vanishes by the algebraic Bianchi identity. Excludes propagating-torsion, fermion-loop, dynamical-Immirzi, non-minimal-matter sectors.
R1 (NJL four-fermion contact)	The torsion-induced contact term sourcing coherent $w = -1$ dark energy	(II)+(III): parity-even structure ($\langle J^5 \rangle \approx 0$, no coherent mean field) is a Tier-II algebraic fact; the M_{Pl}^2 amplitude suppression is a Tier-III order-of-magnitude estimate from a standard Cartan torsion-elimination derivation: the companion’s coefficient-one benchmark $\kappa n_\psi^2 / \rho_\Lambda \simeq 3.6 \times 10^{-69} (n_\psi / 100 \text{ cm}^{-3})^2$, i.e. ≈ 68 orders below ρ_Λ even at a deliberately elevated ISM-like normalization [1].
R2 (one-loop graviton corrections)	One-loop promotion of γ_{BI} reaching the observed birefringence amplitude	(III) Ansatz-level : amplitude budget suppressed by H_0/M_{Pl} and $\alpha_{\text{em}}/(4\pi)$ under an explicitly-labeled scaling ansatz; exploratory, not load-bearing.
R3 (Immirzi running)	Quantum running of γ_{BI} generating a dark-energy-relevant shift	(II)+(III): mass-dimension lock is structural; the chiral-count ansatz bound (motivated by Date–Kaul–Sengupta [8]) is a Tier-III order-of-magnitude upper bound, deliberately loose and non-load-bearing ($\gtrsim 60$ orders of margin); the integrated Benedetti–Speziale flow gives the far smaller derived primary estimate (Sec. IV).
R4 (parity-odd CMB / spectator-ALP coupling)	Minimal ECH <i>deriving</i> the single coupling that yields both β_{obs} and ρ_Λ	(II) Structural naturalness objection , <i>not</i> an amplitude exclusion: a free- α/M ALP fit reproduces β_{obs} , but ECH supplies neither $m_\theta \sim H_0$ nor the fitted α/M , relocating the CC problem. Conditional on the on-shell scaling ansatz (App. A).

Physical foundation: the dimension-4 off-shell operator basis.—The formally dimension-4, gauge-invariant, diffeomorphism-covariant off-shell operator basis is the closed local expansion

$$S_{\text{eff}}^{(4)} = \int d^4x \sqrt{-g} \sum_n c_n \mathcal{O}_n^{[4]}, \quad \{\mathcal{O}_n^{[4]}\} = \{\text{O1–O6}\}, \quad (7)$$

where each $\mathcal{O}_n^{[4]}$ is a genuine dimension-4 Lagrangian density and each Wilson coefficient c_n is a dimensionless rational. The operators $\{\text{O1–O6}\}$ constitute the *primary physical foundation* of the no-go: the single-scale NDA ceiling, the Fierz closure, and the Bianchi-vanishing re-

sults below are all derived properties of this off-shell dimension-4 basis; none depends on the on-shell shorthand Eq. (6) for its logical force. *Construction rule.*—The admitted building blocks are the tetrad e_μ^I (with the invariant tensors ε and η), the curvature two-form of the torsionful connection, the algebraically-fixed torsion $T = \kappa S$, and the minimal axial current J^5 ; densities are formed at zero additional derivative order (no derivatives beyond those internal to R and T), with a parity-odd ε contraction, full index contraction to a scalar density, and mass dimension exactly four after the explicit M_{Pl}^2 promotions of Eq. (8). Under this rule the basis is the six-member set of Eq. (8). Completeness of the list is

asserted from the construction rule — we know of no admitted contraction outside it — and is not established by an exhaustive symbolic enumeration; the released script verifies the two reduction identities used below, not the enumeration itself. One density the zero-derivative rule silently excludes deserves explicit disposal, since it is gauge-invariant, diffeomorphism-covariant, parity-odd, and of dimension exactly four: $\sqrt{-g} \nabla_\mu J^{5\mu}$. It is an exact total derivative, $\sqrt{-g} \nabla_\mu J^{5\mu} = \partial_\mu(\sqrt{-g} J^{5\mu})$, contributing zero to the local equations of motion and zero to the vacuum energy, exactly as O2/O3 below; at the quantum level its anomalous content introduces $F\tilde{F}$ only once electromagnetic fields — outside the minimal field content — are added. The rule-based completeness assertion therefore survives this obvious derivative-order probe. Because the underlying schematic invariants carry different naive dimensions under the paper’s conventions ($[e] = 0$, $[R] = +2$, physical torsion $[T] = [\kappa S] = +1$, spin current $[S] = [J^5] = +3$, $[\kappa] = [M_{\text{Pl}}^{-2}] = -2$), the operator *definitions* include the explicit compensating M_{Pl} power that renders each density dimension-4; we do *not* attach a dimensionless c_n to a lower-dimension invariant. *Shorthand and its relation to the basis.*—Eq. (6) with its dimension-+1 integrand is a phenomenological on-shell representative of the dimension-4 basis in which the coefficient α/M (mass dimension -1) stands in for the basis operators’ M_{Pl}^2 prefactor (mass dimension $+2$) — exactly the three-unit coefficient substitution quantified in Appendix A; it is a convenient dimensional-analysis shorthand, *not* the fundamental action term and *not* load-bearing for the no-go. To clarify the mass-dimension bookkeeping: the off-shell operators O1–O6 are genuinely dimension $+4$ (each carrying the compensating M_{Pl}^2 factor that makes the Wilson coefficient dimensionless), while Eq. (6) represents the same physics on the torsion-free branch with the phenomenological coefficient α/M in place of the basis prefactor M_{Pl}^2 ; the apparent dimension-+1 deficit is entirely that α/M -for- M_{Pl}^2 coefficient substitution (Appendix A), not any ill-defined operator. The Case II on-shell curvature dressing (inserting $R \sim M_{\text{Pl}}^2$ to close the dimension count) is an *illustrative heuristic*, subordinate to the operator basis and used solely for dimensional-analysis aesthetics; it carries no independent logical weight in the no-go argument. Concretely, the six dimension-4 densities are (names as in Table III)

$$\begin{aligned} \mathcal{O}_1^{[4]} &= M_{\text{Pl}}^2 \varepsilon^I e^J R_{IJ}, & \mathcal{O}_2^{[4]} &= M_{\text{Pl}}^2 \text{NY}, \\ \mathcal{O}_3^{[4]} &= R^{IJ} \wedge R_{IJ}, & \mathcal{O}_4^{[4]} &= M_{\text{Pl}}^2 \varepsilon^{\mu\nu\rho\sigma} T^I{}_{\mu\nu} T_{I\rho\sigma}, \\ \mathcal{O}_5^{[4]} &= \varepsilon T e J^5, & \mathcal{O}_6^{[4]} &= M_{\text{Pl}}^2 \varepsilon^{\mu\nu\rho\sigma} R_{\mu\nu\rho\sigma}, \end{aligned} \quad (8)$$

Here the torsion two-form T^I carries a single internal index (components $T^I{}_{\mu\nu}$; the component tensor $T^{abc} = \kappa S^{abc}$ of Sec. II is obtained by converting the spacetime pair with the tetrad), so the torsion-square density O4 is the single-internal-contraction invariant $T_I \wedge T^I$ written in components — the torsion piece of the Nieh–Yan

identity $d(e_I \wedge T^I) = T_I \wedge T^I - e_I \wedge e_J \wedge R^{IJ}$ [16]. The basis is arranged so that Pontryagin (O3, $[R \wedge R] = +4$) and axial-torsion (O5, $[TeJ^5] = +1 + 0 + 3 = +4$) already sit at dimension 4 with a bare dimensionless coefficient, while the Holst dual, Nieh–Yan, torsion-square, and single-curvature densities (each of naive dimension $+2$ as component densities) are promoted to dimension 4 by the single available scale M_{Pl}^2 inside the operator definition. This bookkeeping is essential: labelling the bare invariants as “dimension-4 with dimensionless c_n ” would misstate their mass dimension, whereas the definitions above carry the M_{Pl}^2 prefactors explicitly so that every $c_n \mathcal{O}_n^{[4]}$ is a bona-fide dimension-4 density. The physics conclusion is unchanged: the algebraic Cartan constraint $T = \kappa S$, the first (algebraic) Bianchi identity, and Chern–Weil exactness collapse this set to (i) topological total derivatives (O2, O3), (ii) the M_{Pl}^{-2} -suppressed Fierz-closed four-fermion sector (O4, O5), or (iii) Bianchi-vanishing single-curvature densities (O1, O6).

Main-text closure argument.—The collapse is closed by the exhibited finite basis plus two tensor identities; the identities are verified symbolically in the released script `arxiv/scripts/dim4_parityodd_enumeration.py` (which, its filename notwithstanding, verifies the two identities and performs no basis enumeration; the enumeration rests on the construction rule above). (a) The single-curvature parity-odd densities vanish identically: for any Riemann tensor obeying the pair antisymmetries and the algebraic Bianchi identity $R_{\mu[\nu\rho\sigma]} = 0$, $\varepsilon^{\mu\nu\rho\sigma} R_{\mu\nu\rho\sigma} = 0$, killing both O1 and O6 on the torsion-free branch. (b) The torsion-square and axial-torsion densities reduce under $T = \kappa S$, using $S_{abc} S^{abc} = -\frac{3}{8} (J^5 \cdot J^5)$ (from $S^{abc} = \frac{1}{4} \varepsilon^{abcd} J_d^5$ and the Lorentzian $\varepsilon_{abcd} \varepsilon^{abce} = -3! \delta_d^e$, consistent with the companion paper’s torsion-elimination normalization [1]), to the parity-odd four-fermion contact operator $\kappa^2 (J^5 \cdot J^5)$, which the Fierz-by-Fierz projection lemma (Appendix C) shows closes onto the finite $\{SS, VV, AA, PP\}$ basis with *no* escape operator and *no* change of M_{Pl}^{-2} power. (c) O2 (Nieh–Yan) and O3 (Pontryagin) are exact total derivatives, contributing zero to the local equations of motion and zero to the vacuum energy. Hence every admissible local dimension-4 parity-odd density in minimal ECH is topological, Fierz-basis-reducible, or Bianchi-vanishing, and none produces a (meV)⁴ vacuum energy without a new light scale $\mu \ll M_{\text{Pl}}$ or a protected symmetry — either of which is itself the tuning to be explained. The concrete operator definitions, the symbolic checks, and the Fierz matrix are given in full in Appendices A1–C (Table III). On the on-shell branch (algebraic torsion eliminated; O1/O6 vanishing identically by the algebraic Bianchi identity, O2/O3 total derivatives), this closed dimension-4 set is represented by the single dimension-+1 shorthand of Eq. (6), whose coefficient α/M (dimension -1) replaces the basis operators’ M_{Pl}^2 prefactor (dimension $+2$): the +1-vs-+4 counting is exactly that three-unit

coefficient substitution (Appendix A), not a property of an ill-defined action. The identification $\rho_\Lambda = \Xi M_{\text{Pl}}^4$ is therefore governed by a *single-scale NDA dimensional no-go* that holds identically for the formal dimension-4 basis and its dimension-+1 shorthand (Appendix A1 verifies the closure survives at genuine dimension 4 without the on-shell dressing): the three-unit deficit forces the missing powers to M_{Pl}^3 (the only scale in minimal ECH), so the natural density is $\sim M_{\text{Pl}}^4$, never $(\text{meV})^4$ (Appendix A). This is a dimensional-impossibility argument, not a derived amplitude, and hence not circular; it is treated as such throughout.

VI. DISCUSSION: SCOPE AND BOUNDARIES

Consistent with the evidentiary classification of Sec. III and Table II, we state explicitly what this survey does and does not establish.

What is established. Thirteen distinct mechanism-class constraints, covering seven foundational classes and six observational branches, collectively close the four enumerated minimal-ECH dark-energy channels (R1–R4) under stated assumptions. Two of those channels (R2, R3) are closed here at the amplitude-budget level — Route 2 with ≈ 60 (conservatively ≥ 58) orders of margin against the observed birefringence amplitude, Route 3 with 61–67 orders against the observed dark-energy density; the closures are robust to $\mathcal{O}(1)$ – $\mathcal{O}(10^{10})$ rescalings of their ansatz-level normalizations. The operator-basis argument (Sec. V) extends the governing dimensional bound from one representative operator to the exhibited six-member local dimension-four parity-odd operator basis {O1–O6} admitted by the minimal-coupling field content, with no escape operator identified within that basis; completeness of the basis is asserted from the stated construction rules, not proved by exhaustive symbolic enumeration.

What is not established. This is a *channel-level*, not a fully operator-level, no-go over the entire space of diffeomorphism-invariant gravitational operators: the operator-basis argument is scoped to local, gauge-invariant, mass-dimension-exactly-four densities built from the minimal-ECH field content (tetrad, algebraically-fixed torsion, Holst term, minimally/totally-antisymmetrically coupled fermions). It does not enumerate derivative four-fermion terms suppressed by additional momentum powers, curvature-torsion mixed invariants of higher order, multi-species chiral structures beyond the single-species minimal content, dynamical-Immirzi-field completions, or non-minimal (trace/tensor) torsion irreps — any of these constitutes a genuine escape from the minimal-coupling scope, not a loophole within it. A non-minimal completion introducing a new light scale $\mu \ll M_{\text{Pl}}$, or a protected symmetry zeroing the Planck-scale piece, would itself evade the single-scale bound — but such a scale or symmetry *is* the fine-tuning the mechanism was meant

to explain, so evading the no-go this way relocates rather than solves the cosmological-constant problem.

The barrier catalog is a map of failure modes of mixed evidentiary strength, not thirteen independently decisive theorems: five entries (B5, B6, B7, B10, B13) are general naturalness or classification arguments rather than ECH-specific calculations, and one (B9) is an explicitly heuristic closure conditional on stated equilibrium assumptions that dissipative or particle-producing bounces can violate. Only the perturbation-transparency theorem (B14; stated and proved self-contained in Appendix D) enters the closure table as a Tier-I rigorous result within its stated scope; the Route-1 torsion-elimination derivation (companion-paper content) is likewise first-principles, but the closure it feeds rests on Tier-II+III arguments (Table II). The R2/R3 amplitude estimates are Tier-III ansatz-level bounds: honest to a factor, not first-principles derivations, though robust to their own ansatz uncertainty by tens of orders of magnitude. The R4 closure is a Tier-II naturalness/explanatory-deficit objection rather than an amplitude exclusion, and its numerical inputs (α/M , β_{obs}) are used here only as context for the Route-2 estimate; its own derivation is a companion-paper (or literature) result, not reproduced in this survey.

Finally, every barrier and closure in this catalog is specific to the *minimal* ECH mechanism class. Other bouncing cosmologies — notably quintom scenarios with independent dynamical fields, or bounces with non-minimal torsion couplings — access observational and structural channels outside this catalog’s scope and are not addressed by any result in this paper.

VII. CONCLUSIONS

We have presented a structural no-go survey of minimal Einstein–Cartan–Holst spin-torsion routes to late-time dark energy and bounce phenomenology. Through systematic analysis of 7 foundation classes and 6 observational branches, we catalog 13 distinct mechanism-class constraints (14 historical entries, B8 subsumed by B14) that map the minimal-ECH parameter space and identify which dark-energy pathways require phenomenological assumption or physics beyond the minimal framework. Two of the four enumerated dark-energy channels, Route 2 (one-loop graviton-sector corrections to the Holst term) and Route 3 (quantum running of the Barbero–Immirzi parameter), are closed here at the amplitude-budget level, one anchored in a published one-loop renormalization and closed with ≈ 60 (conservatively ≥ 58) orders of margin against the observed birefringence amplitude, the other bounded between a derived integrated renormalization-group-flow estimate (~ 67 orders) and a deliberately pessimistic chiral-count bound (~ 61 orders) below the observed dark-energy density. An operator-basis argument extends the governing single-scale dimensional bound from one representative parity-odd operator to the exhibited six-member local dimension-four oper-

ator basis {O1–O6} admitted by the minimal-coupling field content under stated construction rules: every member of that basis is topological, Fierz-basis-reducible, or Bianchi-vanishing, with no identified escape operator within the stated minimal scope (evidentiary status as classified in Sec. V).

The catalog’s collective force is deliberately not overstated: it is presented with an explicit evidentiary classification distinguishing rigorous results, structural arguments, and ansatz-level estimates, and its scope is bounded to the minimal-coupling ECH field content. Non-minimal completions — an added light scale or non-minimal torsion coupling — lie explicitly outside this survey and would themselves reintroduce the fine-tuning the closed channels were meant to avoid. Read together with the companion paper’s torsion-elimination and perturbation-transparency results [1], this survey completes the channel-level accounting of minimal-ECH dark-energy routes: all four enumerated channels close, by amplitude suppression (R1–R3) or by naturalness/explanatory deficit (R4), under the stated minimal-coupling assumptions.

Data and Code Availability

The symbolic verification scripts supporting the operator-basis argument (Sec. V, Appendices A–C) are publicly available in the companion papers’ repository, <https://github.com/Hubify-Projects/bigbounce> (the paths below link to the committed files):

```
arxiv/scripts/dim4_parityodd_enumeration.py
arxiv/scripts/fierz_lemma_check.py
research/theory_audit/fierz_adjudication_2026_08_05.py
research/theory_audit/fierz_adjudication_2026_08_05.md
```

The first script verifies the two reduction identities of Appendix A1 (Checks A and D); the second verifies the companion paper’s Fierz convention (see the convention note in Appendix C); the third independently adjudicates the Fierz coefficient identity of Eq. (C2) by explicit 4×4 Dirac-matrix construction and exact Grassmann-algebra derivation, with the fourth file its human-readable report. None of the scripts performs the basis enumeration, which rests on the stated construction rule. These exact files are frozen at immutable repository commit [9b92721d5d7e1](https://github.com/Hubify-Projects/bigbounce/commit/9b92721d5d7e1). The provenance boundary between the two archives is as follows: the companion paper and the companion’s own verification scripts are preserved as an immutable archival

deposit under [doi:10.5281/zenodo.21481838](https://doi.org/10.5281/zenodo.21481838) (CC-BY-4.0, deposited 2026 July 22) [1], whereas the 2026-08-05 adjudication artifacts listed above post-date that deposit and are available at the pinned repository commit only, not within the Zenodo deposit.

ACKNOWLEDGMENTS

This survey builds on the published spin-torsion cosmology of Nikodem Popławski [17], on the published derivations of Simone Mercuri [7] and of Laurent Freidel, Djordje Minic, and Tatsu Takeuchi [6] connecting the Barbero–Immirzi parameter to parity-violating fermion interactions, and on the published one-loop renormalization results of Ilya Shapiro and Poliane Teixeira [2] and of Dario Benedetti and Simone Speziale [3] that anchor the Route-2/Route-3 budgets. None of the named researchers was involved in this work, and no endorsement is implied.

AI-assisted methods.—This work was carried out with an agentic AI research pipeline built on Anthropic Claude (Opus 4 family, 2026 releases), operated under the author’s direction, used for systematic barrier-cataloging, symbolic and dimensional cross-checking, literature triage, and manuscript preparation. All scientific claims, derivations, and numerical results were verified by the author against the committed computational artifacts in the public reproducibility tree, which together with the repository git history constitutes a public audit trail. The author takes sole responsibility for all scientific claims, derivations, numerical results, and bibliographic attributions in this paper. No external funding was received for this research.

Appendix A: Dimensional Status of the Parity-Odd Operator

The parity-odd operator (Eq. 6) has off-shell mass dimension +1, not the +4 required for a local Lagrangian density:

$$[\alpha/M] = -1, \quad [\epsilon^{\mu\nu\rho\sigma} e_\mu^I e_\nu^J \mathcal{F}_{IJ\rho\sigma}] = +2 \\ \implies [\mathcal{L}_{\text{odd}}] = +1. \quad (\text{A1})$$

Rather than an obstacle to be bridged by assumption, this three-unit deficit *is* the physical content of the no-go, once read through effective-field-theory power counting. A relevant (dimension $d < 4$) operator carries, by naive dimensional analysis, a Wilson coefficient of size Λ^{4-d} set by the EFT cutoff [18–20]. In minimal ECH the *only* available scale is $\Lambda \sim M_{\text{Pl}}$: there is no intermediate threshold and no assumed cancellation. The three missing mass powers are therefore *forced* to be M_{Pl}^3 — one cannot conjure a light scale from a theory that has none — and there are exactly two dimensionally-consistent ways to supply them.

¹ A later repository revision clarifies the first script’s descriptive text only; the checks performed and their pass/fail results are identical.

Case I — coefficient dressing (local NDA reading).—The natural coefficient of the relevant operator is $c \sim M_{\text{Pl}}^{4-1} = M_{\text{Pl}}^3$, and the only static (VEV) piece of $\varepsilon e e \mathcal{F}$ that can act as a cosmological constant is the dimension-+1 torsion/topological density T , giving $\rho \sim M_{\text{Pl}}^3 \langle T \rangle$. A coherent cosmological torsion is forbidden today by parity and isotropy ($\langle T \rangle \rightarrow 0$, hence $\rho \rightarrow 0$), while at the bounce $\langle T \rangle \sim M_{\text{Pl}}$ returns $\rho \sim M_{\text{Pl}}^4$. Neither reproduces $(\text{meV})^4$.

Case II — on-shell curvature dressing.—Retaining the small coefficient α/M (dimension -1) and inserting on-shell bounce curvature $R \sim M_{\text{Pl}}^2$ gives

$$\rho_{\Lambda}^{\text{bounce}} \sim (\alpha/M) M_{\text{Pl}}^5 \sim 10^{-2} M_{\text{Pl}}^4, \quad (\text{A2})$$

a bounce-era density that must dilute by $e^{-3N_{\text{tot}}} \sim 10^{-122}$ ($N_{\text{tot}} \approx 94$) to reach $\rho_{\Lambda}^{\text{obs}}$. To fix the bookkeeping once: Eq. (A2) computes the *local on-shell parity-odd pseudo-density* at the bounce under the Case II dressing; the *generic bounce-scale density* entering the 10^{122} hierarchy below is $\rho_{\text{bounce}} \sim M_{\text{Pl}}^4$, and the difference between the two shifts the required e-fold count by $\mathcal{O}(1)$ (the 92-vs-94 spread quantified in the dependency statement below). The $+1 \rightarrow +4$ mass-dimension promotion in Case II is a *heuristic dimensional argument*: absorbing the three missing mass powers into on-shell bounce curvature factors ($R \sim M_{\text{Pl}}^2$) is a dimensional assignment consistent with the single-scale power-counting logic, but it does not constitute a full field-theoretic formalization (e.g., a controlled effective-operator construction demonstrating that the specific curvature insertion is gauge-invariant, diffeomorphism-covariant, and free of operator-mixing contributions at the bounce scale). This limitation is explicitly disclosed; the NDA order-of-magnitude conclusion that the natural density sits at $\rho_{\Lambda}^{\text{ECH}} \sim M_{\text{Pl}}^4$ is unchanged by this caveat, since both Case I and Case II land at the same Planck-scale ceiling via the single-scale power-counting logic, regardless of which off-shell completion is adopted.

The single-scale NDA no-go.—In both admissible completions the three missing mass powers are M_{Pl} (the only scale), so any dimensionally-consistent minimal-ECH parity-odd source sits at

$$\rho_{\Lambda}^{\text{ECH}} \sim M_{\text{Pl}}^4 \text{ (up to } \mathcal{O}(1)\text{--}\mathcal{O}(10^{-2})\text{)}, \quad \text{never } (\text{meV})^4. \quad (\text{A3})$$

Torsion cannot be dark energy because dimensional analysis pins its natural density at M_{Pl}^4 and minimal ECH supplies no light scale (no $m_{\theta} \sim H_0$, no dynamically-small α/M) to bridge the resulting ~ 122 -order hierarchy: the cosmological-constant problem re-appears untouched. This is a single-scale power-counting bound, not a fitted amplitude. Because *no* positive ρ_{Λ} is derived (and none is needed), the closure is a naturalness / amplitude-ceiling no-go and *cannot be circular*: a charge of circularity requires a claimed derivation of the conclusion from an assumption of that same conclusion, and here there is no derived amplitude at all — the $+1 \rightarrow +4$ dimensional gap, read correctly, is itself the mechanism.

Residual assumption (explicit, not fabricated).—The bound assumes single-scale EFT: $\Lambda \sim M_{\text{Pl}}$, no intermediate threshold $\mu \ll M_{\text{Pl}}$, and no exact symmetry/topological cancellation. A non-minimal UV completion introducing a new light scale $\mu \ll M_{\text{Pl}}$ in the coefficient ($c \sim \mu^a M_{\text{Pl}}^{3-a}$), or a symmetry that zeroes the M_{Pl}^4 piece and leaves a protected small remainder, could evade the estimate — but such a scale or symmetry *is* the tuning the mechanism was meant to explain, relocating rather than solving the cosmological-constant problem. Minimal ECH contains neither, which is precisely why the no-go holds *within minimal ECH*. Full closure would require a matching calculation over parity-odd ECH-compatible completions showing that any such completion either (i) reproduces the M_{Pl}^4 NDA estimate, or (ii) requires an added light scale/symmetry that is itself the tuning being explained; that matching calculation is left to a companion treatment. For the minimal-ECH field content this matching is in fact immediate rather than deferred: the finite operator basis fixed by the algebraic Cartan constraint and the totally-antisymmetric minimal spin current (Sec. II) contains no operator of lower off-shell dimension than the +1 operator bounded above, so single-scale NDA applied to that ceiling bounds every member of the basis; only a non-minimal light scale or protected cancellation—case (ii)—evades it, and that is the tuning being explained.

Inflationary dilution ($\mathcal{D}_{\text{inf}} \sim e^{-3N_{\text{tot}}}$, with $\mathcal{D}_{\text{inf}} \equiv e^{-3N_{\text{tot}}}$) then yields $\rho_{\Lambda} = \Xi M_{\text{Pl}}^4$ with $\Xi = (\alpha/M) M_{\text{Pl}} \mathcal{D}_{\text{inf}}$ under Eq. (A2). The genuine cosmological-constant hierarchy, evaluated at the exact values fixed above, is $M_{\text{Pl}}^4 / \rho_{\Lambda}^{\text{obs}} = (1.22 \times 10^{19} \text{ GeV})^4 / (2.25 \times 10^{-3} \text{ eV})^4 \approx 8.7 \times 10^{122} \approx 10^{123}$ (rounding the inputs to bare powers of ten would instead give 10^{124} ; we therefore quote the exact-value figure). Rounding convention: we describe this hierarchy as “ ~ 120 orders of magnitude” and adopt the round number 10^{122} in the dilution bookkeeping below — the 10^{122} -vs- 10^{123} choice shifts the required e-fold count by $\ln 10/3 \approx 0.8$, well inside the quoted 92–94 spread (as fixed above, the density entering this hierarchy is the generic bounce-scale $\rho_{\text{bounce}} \sim M_{\text{Pl}}^4$, while Eq. (A2) computes the local Case-II pseudo-density $\sim 10^{-2} M_{\text{Pl}}^4$). The dilution factor required to bridge $\sim 10^{122}$ from M_{Pl}^4 down to the observed ρ_{Λ} is $\mathcal{D}_{\text{inf}} \sim e^{-3N_{\text{tot}}} \sim 10^{-122}$, giving $N_{\text{tot}} \approx 122 \ln 10/3 \approx 94$ e-folds.

Sharper dependency statement.—The precise value of N_{tot} *does* depend on the ansatz choice in Eq. (A2); however, the *overall scale separation* between $\rho_{\Lambda}^{\text{bounce}}$ and $\rho_{\Lambda}^{\text{obs}}$ is ~ 120 orders of magnitude under *any* non-cancellation assumption (whether the bounce-scale density is M_{Pl}^4 , $10^{-2} M_{\text{Pl}}^4$, or $10^{-4} M_{\text{Pl}}^4$ shifts the required e-fold count by $\mathcal{O}(\text{a few})$, not by orders of magnitude), and this scale separation is what drives the 13 mechanism-class structural barriers (Sec. III). The barriers close the bounce-to-dark-energy route on amplitude-budget and operator-counting grounds whose qualitative conclusions are ansatz-independent; only the precision of the headline

N_{tot} figure is ansatz-dependent. Readers should therefore treat $N_{\text{tot}} \approx 92\text{--}94$ as an order-of-magnitude estimate, not as a precise number; the broader structural closure does not depend on the precise e-fold value.

The on-shell scaling of Case II is labeled explicitly as a phenomenological ansatz rather than as a derivation. The controlled EFT-level construction it stands in for — an *operator-basis closure* for the genuine local dimension-+4 parity-odd sector of minimal ECH — is, however, not deferred: it is supplied in the next subsection, where we exhibit that basis and show that single-scale NDA closure survives without the Case II dressing.

1. Genuine Dimension-Four Parity-Odd Completion: the Operator Basis Survives Single-Scale Closure

The dimension-+1 operator of Eq. (6) is a presentationally-simplified representative. A referee may reasonably ask for the *fully gauge-invariant, diffeomorphism-covariant, local dimension-4* completion directly. We therefore exhibit every rule-admitted local, Lorentz- and diffeomorphism-invariant, parity-odd density of mass dimension exactly +4 of the minimal-ECH field content — tetrad e^I_μ , Holst term, *algebraic* torsion fixed by the Cartan constraint $T^{abc} = \kappa S^{abc}$ ($\kappa \sim M_{\text{Pl}}^{-2}$), and minimally-coupled canonical matter (the construction rule of Sec. V) — and determine the fate of each. Completeness of the resulting six-member list is asserted from that construction rule, not established by an exhaustive symbolic enumeration. No new light scale $\mu \ll M_{\text{Pl}}$, no dynamical Immirzi field, no propagating torsion, and no non-minimal (trace/tensor) torsion irreps are admitted, as these lie outside the stated minimal scope. Building-block dimensions: $[e] = 0$, $[\omega] = +1$, $[R] = +2$, $[S] = +3$ (Dirac bilinear), so via Cartan each torsion factor carries $[\kappa S] = +1$. Under these assignments the bare invariants $\varepsilon e e R$, Nieh–Yan, and T^2 are component densities of naive dimension +2, while Pontryagin $R \wedge R$ and the axial-torsion current structure $\varepsilon T e J^5$ are already dimension +4. To write a genuine dimension-4 Lagrangian with *dimensionless* Wilson coefficients we therefore include the single available scale M_{Pl}^2 explicitly in each operator definition [Eq. (8)], rather than attaching a dimensionless c_n to a lower-dimension object. Column “dim” below records the naive dimension of the *bare* schematic invariant; the tabulated operator $\mathcal{O}_n^{[4]}$ is that invariant multiplied by the prefactor of column “prefactor” so that $[\mathcal{O}_n^{[4]}] = +4$ throughout. The “Fate” column records the reduction of the *bare* invariant (hence its κ -power tracks the bare dimension: O4’s bare torsion-square $\varepsilon T T \rightarrow \kappa^2(J^5 \cdot J^5)$ at dim +2, O5’s bare $\varepsilon T e J^5 \rightarrow \kappa(J^5 \cdot J^5)$ at dim +4). Restoring each prefactor, the two *genuine* dimension-4 densities land on the same member of the Fierz basis at the same power: $\mathcal{O}_4^{[4]} = M_{\text{Pl}}^2 \cdot \kappa^2(J^5 \cdot J^5) = \kappa(J^5 \cdot J^5)$

[using $M_{\text{Pl}}^2 \kappa^2 = \kappa] = \mathcal{O}_5^{[4]}$, both at dimension +4 and both M_{Pl}^{-2} -suppressed — consistent with the “+4 throughout” contract.

Two tensor identities carry the reduction; both are verified symbolically in the released script `arxiv/scripts/dim4_parityodd_enumeration.py` (whose printed output tags [CHECK A] and [CHECK D] label the two identities; the script contains no other checks).

Check A — single-curvature parity-odd invariants vanish.—For any Riemann tensor obeying the pair anti-symmetries and the first (algebraic) Bianchi identity $R_{\mu[\nu\rho\sigma]} = 0$,

$$\varepsilon^{\mu\nu\rho\sigma} R_{\mu\nu\rho\sigma} = 0 \quad (\text{identically}), \quad (\text{A4})$$

verified by imposing the identity as linear constraints on a generic symbolic curvature and contracting. This is the same Bianchi-vanishing that drives the perturbation-transparency result established in the companion paper’s zero-spin-branch analysis [1]; it kills *both* O1 (the Holst dual) and O6 (any single-curvature parity-odd density) on the torsion-free branch. The torsionful piece of R is $\mathcal{O}(\kappa S)$ and merely feeds the O4/O5 four-fermion channel.

Check D — torsion-square collapse.—The minimal (totally antisymmetric) spin current is dual to the axial vector, $S^{abc} = \frac{1}{4}\varepsilon^{abcd}\bar{\psi}\gamma_d\gamma^5\psi = \frac{1}{4}\varepsilon^{abcd}J_d^5$; using the Lorentzian (mostly-plus, $\varepsilon^{0123} = +1$) contraction $\varepsilon_{abcd}\varepsilon^{abce} = -3!\delta_e^d$ (verified symbolically) gives $S_{abc}S^{abc} = \frac{1}{16}\varepsilon_{abcd}\varepsilon^{abce}J^5_dJ^5_e = -\frac{3}{8}(J^5 \cdot J^5)$, consistent with the companion paper’s torsion-elimination normalization [1]. Every ε -contracted torsion-square (O4) therefore reduces under $T = \kappa S$ to $\kappa^2(J^5 \cdot J^5)$, and the linear axial-torsion coupling (O5) to $\kappa(J^5 \cdot J^5)$ — both members of the Fierz-closed basis (App. C), bounded uniformly at M_{Pl}^{-2} with natural coefficient $\sim M_{\text{Pl}}^4$ by single-scale NDA.

Structural entries.—O2 (Nieh–Yan) is an exact 4-form, $d(e_I \wedge T^I) = T_I \wedge T^I - e_I \wedge e_J \wedge R^{IJ}$ [16], and O3 (Pontryagin) is exact by Chern–Weil; both are total derivatives contributing zero to the local equations of motion and zero to the vacuum energy, independently of torsion.

Verdict.—Every admissible local dimension-4 parity-odd density in minimal ECH *within the enumerated set* {O1–O6} at the stated power-counting order (the non-enumerated classes of the Fierz scope caveat below — derivative four-fermion terms, higher-curvature-torsion mixed invariants, multi-species chiral structures, dynamical-Immirzi coefficients, non-minimal torsion irreps — are explicitly excluded) is (i) a topological total-derivative term (O2, O3) contributing nothing to the EOM or vacuum energy; (ii) a four-fermion contact term (O4, O5) that collapses via the algebraic Cartan constraint to the Fierz-closed basis with natural coefficient $\sim M_{\text{Pl}}^4$; or (iii) a single-curvature parity-odd density (O1, O6) that vanishes by Eq. (A4). No genuine local dimension-4 parity-odd operator in minimal ECH produces a $(\text{meV})^4$ vacuum energy without introducing a new light scale $\mu \ll M_{\text{Pl}}$ or a protected symmetry — and either *is* the tuning the mechanism was meant to

TABLE III. Local dimension-4 parity-odd densities of minimal ECH, their coefficient dimensions, and their fate. The bare schematic invariant of column 3 has naive dimension “dim”; the genuine dimension-4 density $\mathcal{O}_n^{[4]}$ [Eq. (8)] is that invariant times the listed prefactor, carried by a *dimensionless* Wilson coefficient c_n . The “Fate (bare)” column records the reduction of the *bare* invariant, prior to multiplication by the prefactor; the “Final” column restores the prefactor, so the table itself shows the two genuine dimension-4 densities O4 and O5 landing on the same operator $\kappa(J^5 \cdot J^5)$ [$M_{\text{Pl}}^2 \kappa^2 = \kappa$; see text]. All entries are topological total derivatives ($\rightarrow$ zero equations of motion), reduce under the algebraic Cartan constraint to the Fierz-closed four-fermion basis (App. C, natural coefficient $\sim M_{\text{Pl}}^4$ by single-scale NDA), or vanish by the first (algebraic) Bianchi identity. None yields a $(\text{meV})^4$ vacuum energy without a new light scale.

#	dim (bare)	Bare invariant (schematic)	Prefactor	Fate (bare)	Final ($\times$ prefactor)
O1	2	$\varepsilon e^I e^J R_{IJ}$ (Holst dual)	M_{Pl}^2	vanishes (Bianchi, Check A)	0
O2	2	$\text{NY} = d(e_I \wedge T^I)$ (Nieh–Yan)	M_{Pl}^2	exact 4-form $\rightarrow$ 0 EOM	0 (EOM)
O3	4	$R^{IJ} \wedge R_{IJ}$ (Pontryagin)	1	exact $\rightarrow$ 0 EOM	0 (EOM)
O4	2	$\varepsilon^{\mu\nu\rho\sigma} T^I{}_{\mu\nu} T_{I\rho\sigma}$ (torsion ²)	M_{Pl}^2	$\rightarrow \kappa^2(J^5 \cdot J^5)$, Fierz basis	$\kappa(J^5 \cdot J^5)$
O5	4	$\varepsilon T e J^5$ (axial-torsion)	1	$\rightarrow \kappa(J^5 \cdot J^5)$, Fierz basis	$\kappa(J^5 \cdot J^5)$
O6	2	$\varepsilon^{\mu\nu\rho\sigma} R_{\mu\nu\rho\sigma}$	M_{Pl}^2	vanishes (Bianchi, Check A)	0

explain. **Single-scale NDA closure therefore survives at genuine dimension 4**, without recourse to the Case II on-shell dressing: the $+1 \rightarrow +4$ dressing is dispensable once the true dimension-4 basis is exhibited, because every element is topological, Fierz-basis-reducible, or Bianchi-vanishing. The dimension-+1 representative of Eq. (6) is a simplified stand-in for this closed set, not an ad-hoc object requiring an uncontrolled promotion.

Appendix B: Parity Classification of the Route-2 Operator

For ϑ_{NY} a pseudoscalar, $\partial_\mu \vartheta_{\text{NY}}$ transforms as a pseudo-co-vector and $J^{5\mu}$ as a pseudo-vector, so their Lorentz-scalar contraction $\partial_\mu \vartheta_{\text{NY}} J^{5\mu}$ is parity-*even* as a Lagrangian term. The parity-violating phenomenology of Route 2 arises not from the operator’s intrinsic P transformation but from a P-breaking *background* expectation $\langle \partial_\mu \vartheta_{\text{NY}} \rangle \neq 0$: the time-dependent cosmological Nieh–Yan pseudoscalar selects a preferred temporal orientation, spontaneously breaking P and T. Throughout Route 2 (section heading, superscript of Eq. (1), and surrounding text) the label “parity-odd” is therefore to be read strictly as *parity-violating in effect* via this P-breaking background — never as a claim that the Lagrangian density is P-odd under an unbroken parity transformation; the label is retained only for continuity with the established Route-2 terminology.

The photon-coupling chain used in the dimensionless ratio Eq. (2) proceeds via the standard chiral-anomaly relation $\partial_\mu J^{5\mu} \supset (\alpha_{\text{em}}/4\pi) F\tilde{F}$ at the EFT level — the operator of Eq. (1) does not itself contain an electromagnetic field strength — and the resulting β estimate is treated strictly as an amplitude-budget bound, not a derived prediction.

Appendix C: Fierz-by-Fierz Projection Lemma for the Minimal-ECH Four-Fermion Basis

The torsion-elimination step (established in the companion paper’s torsion-elimination analysis [1]) generates, at the single prefactor $\kappa = 8\pi G \sim M_{\text{Pl}}^{-2}$, the axial-axial contact term $c_{AA}(J^5 \cdot J^5)$ (the only structure present under *minimal* fermion coupling, per Freidel–Minic–Takeuchi [6]) and—only once a *non-minimal* fermion coupling is admitted—the vector-axial Holst partner $c_{VA}(J \cdot J^5)$. We prove that the complete Fierz rearrangement of these operators closes onto the enumerated dimension- ≤ 6 parity-relevant basis with *no* escape operator and *no* change of M_{Pl} power.

On the 16-dimensional Clifford basis, grouped into the five Lorentz classes $(S, V, T, A, P) = (1, \gamma^\mu, \sigma^{\mu\nu}, \gamma^\mu \gamma^5, \gamma^5)$, we first state the c-number spinor rearrangement in the normalized convention of Itzykson–Zuber and Nieves–Pal [21, 22]: $e_I(1234) = \sum_J (F_c)_{IJ} e_J(1432)$, with rows I the source classes and columns J the produced classes,

$$F_c = \frac{1}{4} \begin{pmatrix} 1 & 1 & \frac{1}{2} & -1 & 1 \\ 4 & -2 & 0 & -2 & -4 \\ 12 & 0 & -2 & 0 & 12 \\ -4 & -2 & 0 & -2 & 4 \\ 1 & -1 & \frac{1}{2} & 1 & 1 \end{pmatrix}, \quad F_c^2 = \mathbb{1}, \quad (\text{C1})$$

in the order (S, V, T, A, P) . In this normalized basis $\Gamma_A^\mu = i\gamma^\mu \gamma^5$ with $n_A = -i$, and the two factors combine to make e_A exactly the physical $(\bar{\psi} \gamma^\mu \gamma^5 \psi)(\bar{\psi} \gamma_\mu \gamma^5 \psi)$ operator; the phases are already absorbed in F_c . Turning the c-number identity into an anticommuting field-operator identity requires one Grassmann exchange, $F_{\text{op}} = -F_c$. The axial source is the fourth row of Eq. (C1), $\frac{1}{4}(-4, -2, 0, -2, 4) = (-1, -\frac{1}{2}, 0, -\frac{1}{2}, 1)$, so $F_{\text{op}} = -F_c$ yields the operator row $(1, \frac{1}{2}, 0, \frac{1}{2}, -1)$ and the operator-level decomposition

$$(J^5 \cdot J^5) \xrightarrow{\text{Fierz}} SS + \frac{1}{2}VV + \frac{1}{2}AA - PP, \quad (\text{C2})$$

while the parity-odd ($J \cdot J^5$) Holst partner rotates only within the $\{V, A\}$ block (operator-level cross coefficients $F_{VA}^{\text{op}} = F_{AV}^{\text{op}} = \frac{1}{2}$), remaining a dimension-6 parity-odd cross term. The scalar bridge used by the companion paper’s mean-field analysis is visible in the same chain: $(F_c)_{AS} = \frac{1}{4}(-4) = -1$, hence $(F_{\text{op}})_{AS} = +1$ and $G_s = (-3\kappa/16)(+1) = -3\kappa/16$. *Convention note.*—Individual Fierz coefficients are convention-dependent (class normalization, row-versus-column orientation, and the c-number-versus-anticommuting-operator map); Eq. (C2) is stated in the normalized Nieves–Pal c-number basis with the explicit Grassmann-exchange step, matching the companion paper’s published Fierz appendix [1] and giving the γ -independent mean-field scalar-channel factor $G_s = -3\kappa/16$ of the companion’s gap-equation convention. The Fierz rearrangement supplies only this channel factor; the full Sec. II contact operator carries the additional finite-Holst prefactor $\gamma^2/(1 + \gamma^2) < 1$, which arises from the Cartan torsion elimination (not from the Fierz step) and which the companion’s gap-equation bound conservatively omits because it can only reduce the coupling (equivalently, $G_s = -3\kappa/16$ is the $\gamma \rightarrow \infty$ limit) [1]. This coefficient set is independently verified by explicit 4×4 Dirac-matrix construction in both metric signatures and an exact Grassmann-algebra derivation of the operator row (the unique solution for identical fields); see the released verification artifact and its report listed in the Data and Code Availability statement. The released script `arxiv/scripts/fierz_lemma_check.py` independently checks the same convention’s matrix involution and axial-row coefficients. None of this survey’s conclusions depends on the convention-bound coefficients beyond this identity: the produced classes close on $\{SS, VV, AA, PP\}$ (the tensor channel is absent from the axial row), every coefficient is a dimensionless $O(1)$ rational, and the $\kappa \sim M_{\text{Pl}}^{-2}$ prefactor is preserved exactly — which is all the projection lemma uses. Every produced structure lies in the closed set $\{SS, PP, VV, AA\}$ (with $\{V, A\}$ for the cross term); the tensor channel does not appear, and no operator escapes the enumerated basis. Because every coefficient in F is a dimensionless rational, the Fierz map is an $O(1)$ recombination of the *same* four ψ fields and therefore preserves the $\kappa \sim M_{\text{Pl}}^{-2}$ prefactor exactly. This is the explicit Fierz-by-Fierz projection lemma: within single-species minimal ECH the generated four-fermion tower is Fierz-closed and uniformly M_{Pl}^{-2} -suppressed, so the single-scale NDA ceiling of App. A that bounds one representative bounds the entire finite basis. (For multiple fermion species the Lorentz-class closure and the M_{Pl} -power argument are unchanged—Fierz acts on Dirac structure, not flavor—and only the flavor-off-diagonal coefficient bookkeeping is added.) The only operators evading this bound require a *non-minimal* completion (a new light scale $\mu \ll M_{\text{Pl}}$, trace/tensor torsion irreps, or a dynamical Chern–Simons field), which is the stated scope boundary of the no-go.

Scope of the lemma (what it establishes and what it does not).—The Fierz lemma establishes basis-

completeness of the *parity-odd four-fermion contact sector* at M_{Pl} -power-counting within the minimal ECH field content (tetrad, algebraic torsion via Cartan constraint, Holst term, and minimally-coupled fermion). It does *not* enumerate: derivative four-fermion terms suppressed by additional powers of momentum/ M_{Pl} ; curvature-torsion mixed invariants built from higher-order contractions of the Riemann tensor and spin current; chiral/flavor structures involving independent fermion species beyond the single-species minimal content; dynamical topological coefficients arising from a promoted Immirzi field; or non-minimal torsion irreps (trace-vector and tensor irreps) that appear only when the minimal coupling assumption is relaxed. The claim is scoped to the minimal-ECH field content at the stated power-counting order; it is not a complete operator-level no-go over all diffeomorphism-invariant operators in the full gravitational EFT space.

Appendix D: Self-Contained Statement and Proof of the Perturbation-Transparency Theorem (B14)

So that the catalog’s sole Tier-I leg can be refereed from this manuscript, we carry here, faithfully from the companion paper’s zero-spin-branch analysis [1], the statement and proof of the perturbation-transparency theorem. The tensor-sector extension, the explicit second-order Holst-term verification, and the discussion of what would break the transparency are given in full in the companion; nothing below is new to this survey.

Statement.—Consider the classical ECH action with an invertible tetrad, canonical scalar matter, and a real finite nonzero constant γ , on a local patch with matched background, initial, and boundary data (standard falloff conditions, boundary contribution to the first-order variation set to zero). Solve the algebraic connection equation before expanding in perturbations. Then the reduced action on the zero-spin branch is the Einstein–scalar action, so scalar equations and tensor evolution operators coincide with those of general relativity at every perturbative order for the same background and boundary data. Equality of right- and left-helicity *solutions*, rather than of their evolution operators, additionally requires parity-symmetric initial data. The complex singular values $\gamma = \pm i$, nontrivial global or topological sectors, nonstandard boundary terms, quantum loops or anomalies, fermion sources, non-minimal matter, a dynamical Immirzi field, and propagating torsion are outside this statement. This is equality *after* solving the Cartan equation — an on-connection-shell equality of local classical reduced actions — not an off-shell equality of the original first-order actions; its all-orders content does not come from an order-by-order calculation, because the Cartan constraint is algebraic and the Bianchi identity used below is pointwise.

Proof.—(1) *Zero spin density.* A canonical scalar field has zero spin density. (2) *Zero torsion.* The Holst-modified connection equation acts through an algebraic

bivector operator whose inverse exists for real finite nonzero γ (because $1 + \gamma^2 > 0$); the zero scalar source therefore gives $e^{[I} \wedge T^{J]} = 0$, whose invertible-tetrad kernel is trivial: $T^I = 0$. This holds at every classical field configuration in the stated local domain, hence before any perturbative expansion. (3) *The connection reduces to Levi-Civita*, $\Gamma^\lambda_{\mu\nu} = \overset{\circ}{\Gamma}^\lambda_{\mu\nu}$. (4) *The Holst term vanishes by the first Bianchi identity*. Evaluated on the Levi-Civita connection the Holst term reduces to $\frac{1}{2}\epsilon^{\mu\nu\rho\sigma}R_{\mu\nu\rho\sigma}(\overset{\circ}{\Gamma})$, which vanishes identically by the algebraic Bianchi identity $R_{\mu[\nu\rho\sigma]} = 0$ — the same pointwise identity verified symbolically as Check A [Eq. (A4)].

The Holst dual contraction is therefore identically zero on the torsion-free branch (not merely a boundary term) and contributes nothing to the action at any order. No total-derivative argument is a load-bearing step in this proof; the operative statement is the pointwise algebraic Bianchi identity in Step 4. ■

Consequences carried into the catalog: the Holst sector and the Barbero–Immirzi parameter decouple from all scalar and tensor perturbation observables for canonical scalar matter (B14, Branch H, constraining R1–R4), and in particular minimal ECH produces no helicity-dependent tensor propagation from parity-symmetric initial data — the observational content whose channel-level consequence B8 records.

-
- [1] H. Golden, Algebraic Cartan Elimination in Minimal Einstein–Cartan–Holst Gravity: Spin-Sourced Contact and Zero-Spin Scalar Branches, Zenodo [10.5281/zenodo.21481838](https://zenodo.org/record/21481838) (2026), Companion paper (Paper I A), publicly archived on Zenodo under version DOI 10.5281/zenodo.21481838 (concept DOI 10.5281/zenodo.21481837), deposited 2026 July 22 under CC-BY-4.0. Public permanent archive of the manuscript and its exact source; not an arXiv preprint and not peer reviewed.
 - [2] I. L. Shapiro and P. M. Teixeira, Quantum Einstein–Cartan theory with the Holst term, *Classical and Quantum Gravity* **31**, 185002 (2014), [arXiv:1402.4854 \[gr-qc\]](https://arxiv.org/abs/1402.4854).
 - [3] D. Benedetti and S. Speziale, Perturbative running of the Immirzi parameter, *J. Phys. Conf. Ser.* **360**, 012011 (2012), [arXiv:1111.0884 \[hep-th\]](https://arxiv.org/abs/1111.0884).
 - [4] A. Ashtekar and P. Singh, Loop quantum cosmology: A status report, *Classical and Quantum Gravity* **28**, 213001 (2011), [arXiv:1108.0893 \[gr-qc\]](https://arxiv.org/abs/1108.0893).
 - [5] S. Holst, Barbero’s Hamiltonian derived from a generalized Hilbert–Palatini action, *Physical Review D* **53**, 5966 (1996), [arXiv:gr-qc/9511026 \[gr-qc\]](https://arxiv.org/abs/gr-qc/9511026).
 - [6] L. Freidel, D. Minic, and T. Takeuchi, Quantum gravity, torsion, parity violation and all that, *Physical Review D* **72**, 104002 (2005), [arXiv:hep-th/0507253 \[hep-th\]](https://arxiv.org/abs/hep-th/0507253).
 - [7] S. Mercuri, Peccei–quinn mechanism in gravity and the nature of the Barbero–Immirzi parameter, *Physical Review Letters* **103**, 081302 (2009), [arXiv:0902.2764 \[gr-qc\]](https://arxiv.org/abs/0902.2764).
 - [8] G. Date, R. K. Kaul, and S. Sengupta, Topological interpretation of Barbero–Immirzi parameter, *Phys. Rev. D* **79**, 044008 (2009), [arXiv:0811.4496 \[gr-qc\]](https://arxiv.org/abs/0811.4496).
 - [9] D. Benedetti and S. Speziale, Perturbative quantum gravity with the Immirzi parameter, *JHEP* **2011** (6), 107, [arXiv:1104.4028 \[hep-th\]](https://arxiv.org/abs/1104.4028).
 - [10] Y. Minami and E. Komatsu, New extraction of the cosmic birefringence from the Planck 2018 polarization data, *Physical Review Letters* **125**, 221301 (2020), [arXiv:2011.11254 \[astro-ph.CO\]](https://arxiv.org/abs/2011.11254).
 - [11] J. R. Eskilt and E. Komatsu, Improved constraints on cosmic birefringence from the WMAP and Planck cosmic microwave background polarization data, *Phys. Rev. D* **106**, 063503 (2022), [arXiv:2205.13962 \[astro-ph.CO\]](https://arxiv.org/abs/2205.13962).
 - [12] P. Diego-Palazuelos and E. Komatsu, Cosmic birefringence from the Atacama Cosmology Telescope data release 6, (2025), [arXiv:2509.13654 \[astro-ph.CO\]](https://arxiv.org/abs/2509.13654).
 - [13] H. Golden, namaster-proof: Exact pseudo- C_ℓ window inference and content-bound validation for reproducible spin-2 analyses, Zenodo [10.5281/zenodo.21481842](https://zenodo.org/record/21481842) (2026), Companion paper (Paper I B), publicly archived on Zenodo under version DOI 10.5281/zenodo.21481842 (concept DOI 10.5281/zenodo.21481841), deposited 2026 July 22 under CC-BY-4.0. Public permanent archive of the manuscript and its exact source; not an arXiv preprint and not peer reviewed. The software it describes is archived separately at 10.5281/zenodo.21481753.
 - [14] S. M. Carroll, Quintessence and the rest of the world: Suppressing long-range interactions, *Physical Review Letters* **81**, 3067 (1998), [arXiv:astro-ph/9806099 \[astro-ph\]](https://arxiv.org/abs/astro-ph/9806099).
 - [15] Y.-F. Cai, E. N. Saridakis, M. R. Setare, and J.-Q. Xia, Quintom Cosmology: Theoretical implications and observations, *Phys. Rept.* **493**, 1 (2010), [arXiv:0909.2776 \[hep-th\]](https://arxiv.org/abs/0909.2776).
 - [16] H. T. Nieh and M. L. Yan, An identity in riemann–cartan geometry, *J. Math. Phys.* **23**, 373 (1982).
 - [17] N. J. Poplawski, Cosmology with torsion: An alternative to cosmic inflation, *Physics Letters B* **694**, 181 (2010), [arXiv:1007.0587 \[astro-ph.CO\]](https://arxiv.org/abs/1007.0587).
 - [18] G. Buchalla, O. Catà, and C. Krause, On the Power Counting in Effective Field Theories, *Phys. Lett. B* **731**, 80 (2014), [arXiv:1312.5624 \[hep-ph\]](https://arxiv.org/abs/1312.5624).
 - [19] I. Brivio and M. Trott, The Standard Model as an Effective Field Theory, *Phys. Rept.* **793**, 1 (2019), [arXiv:1706.08945 \[hep-ph\]](https://arxiv.org/abs/1706.08945).
 - [20] G. Isidori, F. Wilsch, and D. Wyler, The Standard Model effective field theory at work, *Rev. Mod. Phys.* **96**, 015006 (2024), [arXiv:2303.16922 \[hep-ph\]](https://arxiv.org/abs/2303.16922).
 - [21] C. Itzykson and J.-B. Zuber, *Quantum Field Theory* (McGraw-Hill, 1980).
 - [22] J. F. Nieves and P. B. Pal, Generalized fierz identities, *Am. J. Phys.* **72**, 1100 (2004), [arXiv:hep-ph/0306087 \[hep-ph\]](https://arxiv.org/abs/hep-ph/0306087).