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*

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 fourteen distinct mechanism-class constraints — one per catalog entry; the entries are not restatements of one another, and for the pairs most at risk of subsumption (B8/B14 and B11/B13) the independence is argued explicitly in the text — spanning seven foundational mechanism classes and six observational-channel branches (grouped into four constrained observational channels, since L/M and N/O are treated in pairs and neither N nor O carries a dedicated constraint of its own), which *jointly* close the four routes by which the ECH bounce could plausibly source a Λ-like late-time density. The entries bear on the routes with differing individual force, and two of them (B9, B14) are explicitly *not* used as stand-alone closures of any route; the closure claimed here is that of the catalog taken together, at channel-amplitude granularity. Of the four parity-odd or 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 $\sim$67-order endpoint from the derived integrated flow, the $\sim$61-order endpoint from a deliberately pessimistic chiral-count bound). The flow itself supplies $\Delta\gamma/\gamma$; the step from that to a ρ_Λ ratio is an explicitly labeled Tier-III mass-dimension scaling relation, not a derived stress-tensor matching. Neither route is presented as a viable mapping; both are closed as bounded amplitude budgets. Route 2's *dark-energy* leg, as distinct from its birefringence leg, is closed within the minimal field content by the operator list below and, on the non-minimal dynamical-pseudoscalar branch that its phenomenological representative operator actually requires, only at the naturalness level. We further present an operator-list argument. The object is a six-member *spanning list* {O1–O6} of local, gauge-invariant, diffeomorphism-covariant parity-odd densities of mass dimension exactly four, *algebraic* in the sense of the zero-additional-derivative construction rule stated in the text and exhibited under those construction rules for the minimal-coupling ECH field content. They are parity-odd as ε-contracted densities before on-shell reduction; the surviving members reduce to parity-*even* contact terms. The list is shown to close: every member is a topological total derivative, a Fierz-closed four-fermion contact term suppressed by $M_{\rm Pl}^{-2}$, or identically vanishing. The single-curvature pair vanishes by the algebraic Bianchi identity on the torsion-free branch; on the on-shell branch it reduces to an exact total derivative plus a Fierz-closed contact term, and the ε-contracted torsion-square joins that same contact class, the on-shell Einstein–Cartan–Holst torsion carrying both an axial and a trace-vector irrep (the tensor irrep vanishes identically). The list is deliberately redundant: independent symbolic computation gives it rank four (two exact relations, O1 = O6 and O1 = $\frac{1}{2}$O4 − O2), and rank two modulo total derivatives, so it is a generating set of recognizable invariants and not a linearly independent basis. That the list *spans* the rule-admitted operator space 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 it is Tier-I strictly on the zero-spin (canonical-scalar-matter) branch its hypotheses cover; 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 ($\kappa = 8\pi G \sim M_{\rm Pl}^{-2}$), 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,

* houston@hubify.com

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 $\rho_{\Lambda,\mathrm{obs}} \approx (2.25\,\mathrm{meV})^4$.

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, labeled argument, in some cases by more than one catalog entry acting together: an amplitude-suppression bound, a naturalness/fine-tuning objection, a decoupling identity, or a classification argument. The resulting fourteen-entry catalog yields fourteen distinct mechanism-class constraints 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-list 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 spanning list admitted by the minimal-coupling field content under stated construction rules.

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 (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 rederiving them, and the short derivations a referee needs in order to verify this survey's quantitative claims from this manuscript alone — the torsion-elimination chain fixing the $-(3\kappa/16)[\gamma^2/(1+\gamma^2)]$ contact coefficient and the R1 finite-density benchmark arithmetic — are additionally carried self-contained in Appendix E, following the same faithful-extraction convention as Appendix D. The companion results that remain cited-only (the mean-field NJL gap-equation analysis, the tensor-sector extension of the transparency theorem, and the R4 spectator-ALP consistency check) are used here only with explicit not-peer-reviewed labeling, and their refereeing is part of the companion's own publication path rather than of this survey's. 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-list argument (Sec. V) with its supporting appendices. Route 2 and Route 3 enter as historical candidate mappings, not as results of either paper: the underlying one-loop calculations [2, 3] do not derive the cosmological stress tensor and observable matching a dark-energy mapping would need, and Sec. IV closes them as bounded amplitude budgets without supplying that derivation. This is stated once here and is not repeated at each route.

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-list closure of Sec. V / Appendix A 1 — the six-member genuine dimension-4 parity-odd spanning list {O1–O6} with its script-verified reduction identities and its two exact rank-reducing relations; (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-list 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]; the two that are load-bearing for the identities used below are stated here outright, so that every coefficient in this manuscript can be checked from this manuscript: the metric signature is mostly-plus, $\eta_{ab} = \mathrm{diag}(-,+,+,+)$, and the spacetime Levi-Civita *tensor* is normalized to $\varepsilon^{0123} = +1$, equivalently $\varepsilon_{0123} = -1$ (whence $\varepsilon_{abcd}\varepsilon^{abce} = -3!\,\delta^e_d$, which fixes the $-3/8$ and -3 coefficients of Sec. V and Appendix A 1). We use natural units and $\kappa \equiv 8\pi G$ exactly; in terms of the *reduced* Planck mass $\overline{M}_{\mathrm{Pl}} \equiv (8\pi G)^{-1/2} \simeq 2.44 \times 10^{18}\,\mathrm{GeV}$ this is $\kappa = \overline{M}_{\mathrm{Pl}}^{-2}$, while order-of-magnitude prose throughout writes $\kappa \sim M_{\mathrm{Pl}}^{-2}$ with the *full* Planck mass $M_{\mathrm{Pl}} \equiv G^{-1/2}$ — a deliberate factor-8π abuse (exactly, $\kappa = 8\pi\, M_{\mathrm{Pl}}^{-2}$), never a definition.

Three convention flags are fixed here once. First, the imported one-loop results of Shapiro–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, $\kappa^2 = M_{\rm Pl}^{-4}$ 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 $M_{\rm Pl} \simeq 1.22 \times 10^{19}\,\mathrm{GeV}$; the reduced-versus-full distinction (2.44×10^{18} vs $1.22 \times 10^{19}\,\mathrm{GeV}$) 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 $\kappa n_\psi^2/\rho_\Lambda \simeq 3.6\times10^{-69}$ 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 $\rho_\Lambda \approx (2.3\,\mathrm{meV})^4 \approx 2.8\times10^{-11}\,\mathrm{eV}^4$; under the $\rho_\Lambda = (2.25\,\mathrm{meV})^4$ 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 ($\approx$68 orders below ρ_Λ either way), while the Appendix A hierarchy $M_{\rm Pl}^4/\rho_{\Lambda,\rm obs}$ 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. (7)], M denotes the LQG area-gap scale $M_{\text{area-gap}} \sim M_{\rm Pl}/\sqrt{\gamma}$, 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 $\alpha/M \sim 10^{-21}\,\mathrm{GeV}^{-1}$ (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}. For minimally (totally antisymmetrically) coupled Dirac matter, the spin current is dual to the axial current, $S^{abc} = \frac{1}{4}\varepsilon^{abcd}J_d^5$ with $J^{5\mu} = \bar{\psi}\gamma^\mu\gamma^5\psi$, and solving that constraint — with no restriction on which irreducible parts may appear — returns a torsion with *two* nonvanishing irreps,

$$T_{abc} = \alpha\,\varepsilon_{abcd}J^{5d} + \beta\left(\eta_{ab}J_c^5 - \eta_{ac}J_b^5\right), \qquad \frac{\beta}{\alpha} = \frac{1}{2\gamma}, \tag{1}$$

with $\alpha = \kappa\gamma^2/[2(1+\gamma^2)]$ and $\beta = \kappa\gamma/[4(1+\gamma^2)]$ in the normalization of Eq. (E2), which is the torsion normalization used throughout this survey; the overall torsion sign is a convention, taken positive in the axial channel, and the Holst sign convention flips it and changes nothing else. The axial (4) and trace-vector (4) irreps are both nonzero at every finite nonzero γ, the trace-vector piece being sourced by the axial current through the Holst term; the tensor (16) irrep vanishes identically. Purely axial torsion is the $\gamma \to \infty$ Einstein–Cartan limit, in which $\beta/\alpha \to 0$; at the LQG values $\gamma = 0.2375$ and $\gamma = 0.274$ the trace-vector coefficient is respectively 2.11 and 1.82 times the axial one, so the non-axial piece is the larger of the two. 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]. The elimination chain and the resulting coefficient are carried compactly, faithful to the companion, in Appendix E; this survey begins from that 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 *14 distinct mechanism-class constraints*, one per catalog entry. 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 fourteen 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. That Tier-I status is moreover *branch-scoped*: Appendix D proves the theorem for canonical scalar matter, i.e. on the zero-spin branch, and expressly excludes fermion sources, quantum loops and anomalies, non-minimal matter, a dynamical Immirzi field, and propagating torsion. Nothing outside that branch inherits Tier-I standing. 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, which carries the two logically independent entries B8 and B14) for the possibility of connecting the ECH bounce to late-time dark energy. Each test yielded a named structural constraint. The branch-to-entry multiplicity, stated once so the count is legible in one place: the seven foundations carry one entry each (B1–B7); Branch H carries two (B8, B14); Branch J carries one (B9); Branches L/M carry three between them (B10–B12); and Branches N/O carry one between them (B13). That is $7+2+1+3+1 = 14$ entries across $7+6 = 13$ tested classes and channels — the two grouped branch pairs (L/M and N/O) are treated as single observational channels, and neither N nor O carries a dedicated constraint of its own. 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 $g_{\rm eff} \sim 1/(M_{\rm Pl}\sqrt{|t_3|}) \sim H_0/M_{\rm Pl} \sim 10^{-61}$, 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 $g_{\rm eff} \sim 1$ needs $m_T \sim M_{\rm Pl}$; keeping $m_T \sim H_0$ against a radiative contribution $\delta m_T^2 \sim M_{\rm Pl}^2$ instead requires a cancellation to one part in $(M_{\rm Pl}/H_0)^2 \sim 10^{122}$ — a residual $m_T^2/\delta m_T^2 \sim (H_0/M_{\rm Pl})^2 \sim 10^{-122}$ — 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 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 $m_\phi^2/M_{\rm Pl}^2$ or $(\partial\phi)^2/M_{\rm Pl}^4$; 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 $N_{\rm tot} \approx$ 92–94 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

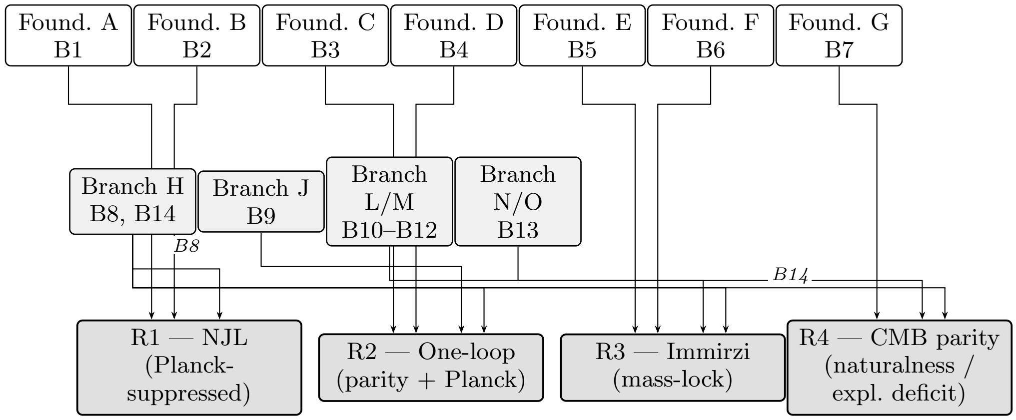

FIG. 1. **Structure of the 14-entry barrier closure.** Top row: the 7 foundation mechanism classes (Foundations A–G; Barriers 1–7, mechanism-class barriers). Middle row: the 6 observational-channel branches (Branches H, J, L, M, N, O; Barriers 8–14, observational-channel barriers), grouped into four 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 14 distinct mechanism-class constraints. Arrows indicate which barrier classes constrain each route. Branch H carries arrows to all four routes; the italic in-figure edge labels give its per-barrier attribution directly: the upper Branch-H arrow to R1 is driven by B8 (label "B8"), the fermionic parity argument, while the lower fan to R2, R3, and R4 is driven by B14 (perturbation transparency, label "B14"), which removes the *classical zero-spin* perturbative baseline of those three routes. B8 and B14 are logically independent: B8 is a statement about the fermionic $(J^5)^2$ sector, whereas B14's hypotheses require zero spin density and expressly exclude fermion sources, so neither implies the other and both are counted. Several entries share the scaling ansatz but each probes a distinct physical failure mode. The diagram records channel structure only, not evidentiary weight: the entries differ in evidentiary status — the sole Tier-I rigorous theorem is B14's perturbation transparency, and only on its zero-spin branch (Appendix D); the remainder are structural or ansatz-level arguments of mixed individual strength, and five entries (B5–B7, B10, B13) are general naturalness or classification arguments — with the per-route classification recorded in Table II.

TABLE I. The 14 catalogue entries — 14 distinct mechanism-class constraints — 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) on *disjoint matter branches*, and are therefore logically independent and separately counted: B8 argues from the parity of the fermionic $(J^5)^2$ contact interaction, while B14's hypotheses require zero spin density (canonical scalar matter) and expressly exclude fermion sources (Appendix D). Neither is a logical consequence of the other, so the criterion of Sec. III counts both.

#	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

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]. 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. B8 is the catalog's constraint on the *fermionic* branch of this observable channel, and it is logically independent of B14: B14's hypotheses require zero spin density, which is precisely the condition the nonzero J^5 of B8 violates, so B14 neither implies nor confirms B8 and the two are counted separately (Table I).

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. The [R2] tag records which route this removes an escape from rather than which amplitude it bounds: R2's one-loop Holst-sector correction is a *small* shift, so the only way it could reach ρ_Λ is if the bounce additionally selected a vacuum on which that shift is amplified or frozen in; B9 removes exactly that selection step, as B3 and B4 remove the amplitude and coupling steps. It is not an independent bound on the one-loop amplitude. 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 $\Omega_{\rm GW}^{\rm ECH}|_{\rm bounce} \lesssim (\rho_{\rm crit}/\rho_{\rm Pl})^2 \simeq 0.07$–$0.17$, using the LQG-bounce critical-density window $\rho_{\rm crit}/\rho_{\rm Pl} \simeq 0.27$–$0.41$: 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 $\gamma \approx 0.274$ (the improved state-counting value of Ghosh & Mitra [5]) into the same $\rho_{\rm crit} = \sqrt{3}/(32\pi^2\gamma^3)\,\rho_{\rm Pl}$ 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 $\Omega_{\rm GW}(f_{\rm nHz}) \sim 10^{-9}$ (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.) B13 is logically independent of B11 under the criterion of Sec. III ("no barrier is a logical consequence of another"): B11 states that the *gauge* sector decouples universally, which by itself says nothing about relative couplings *among fermion species*, and B13 states that the *matter* sector couples democratically, which by itself permits a gauge sector with species-blind but channel-selective couplings. Neither implies the other, and neither is implied by B4 (Planck Suppression), which fixes the *magnitude* of the coupling both invoke but not its *universality* across sectors. They are counted separately for that reason.

B14 — Perturbation Transparency (Branch H, ECH perturbation gates). *[R2–R4, zero-spin branch.]* 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, scalar and tensor sectors alike; the explicit order-by-order second-order Holst-term verification is elaborated in the companion paper's zero-spin-branch analysis [1]. This is the catalog's sole Tier-I closure leg (Table II), and its route reach is fixed by its hypotheses rather than asserted. Because the theorem assumes zero spin density — and explicitly excludes fermion sources, quantum loops and anomalies, non-minimal matter, a dynamical Immirzi field, and propagating torsion — it does *not* bear on R1, whose defining object is the nonzero fermionic axial current (R1 is constrained instead by B1, B2 and B8). What B14 does establish for R2, R3 and R4 is that their *classical zero-spin* perturbative baseline is exactly inert: minimal ECH reproduces the Einstein–scalar perturbation sector at every order, so any signal in those routes must come entirely from the quantum or non-minimal ingredient that defines them — which is what the amplitude budgets of Sec. IV and the naturalness argument of Sec. IV C then bound. B14 is not, and is not used as, a closure of the fermionic or one-loop content of any route.

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

The two routes closed in this section are the *historical* candidate mappings of Sec. I — one-loop Holst-sector corrections (Route 2) and Immirzi-parameter running (Route 3) mapped onto a late-time dark-energy density. 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. (3)], for Route 3 against the *observed dark-energy density.* The published one-loop inputs enter as normalization anchors for the budgets, not as derivations of a cosmological observable. We state the division of labour once, at the head of the section, because the two routes are closed against different observables and, for Route 2, against different observables by different arguments. Every quantitative object in Sec. IV A is a CMB rotation angle: what is computed there is the closure of Route 2's *birefringence* channel. Route 2's *dark-energy* leg is not computed in that subsection and is not carried by that budget. It splits on whether the Nieh–Yan pseudoscalar coefficient is constant: *(i)* within the minimal field content the Immirzi parameter, and with it the Nieh–Yan coefficient, is a *constant* fixed by the LQG area spectrum (B7), and the corresponding densities are the Holst dual and Nieh–Yan entries O1/O2 of the Sec. V list — O2 an exact total derivative, and O1 on shell that same total derivative plus a Fierz-closed contact term suppressed by $M_{\rm Pl}^{-2}$ [Eq. (13)] — so minimal Route 2 sources no dark energy beyond a Planck-suppressed contact term already inside the bounded four-fermion sector, by the operator-list argument; *(ii)* the phenomenological representative Eq. (2) used for the birefringence budget is *not* in that list — it promotes the coefficient to a dynamical light pseudoscalar $\vartheta_{\rm NY}(x)$ and is a dimension-five operator carried by a dimension-(-1) coefficient — so its dark-energy leg is *not* closed by the single-scale NDA bound, which App. A states explicitly can be evaded by exactly such a light scale. On that branch Route 2's dark-energy leg is closed only at the naturalness level, as a Tier-II R4-class objection reinforced by B7; this is set out at the end of Sec. IV A. Route 3, by contrast, is bounded directly against $\rho_{\Lambda,\rm obs}$ in Sec. IV B.

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 on shell, once torsion is integrated out, to the Nieh–Yan density *plus* the ε-contracted torsion-square of Eq. (13) [6, 7] — the latter vanishing only in the $\gamma \to \infty$ Einstein–Cartan limit, where the trace-vector torsion irrep of Eq. (1) switches off. 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]; the classical Holst-plus-fermion and Nieh–Yan structure on which that analysis rests is due to Mercuri [8], which is a classical construction and not itself a renormalization-group computation. 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

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

where $\vartheta_{\rm NY}(x)$ is the Nieh–Yan pseudoscalar (mass dimension $+1$; written $\vartheta_{\rm NY}$ 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 the slowly varying $\mathcal{O}(1)$ renormalization-group function of γ supplied by the Shapiro–Teixeira one-loop α_4/Ω_{4x} family — fixed below, at Eq. (3)'s discussion, as $|\Omega_{44}/\alpha_4| \approx 3.3$ at the LQG value $\gamma \approx 0.24$, so that no property of $\beta(\gamma)$ used here is left to a forward reference. The parity classification of Eq. (2) — 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. (2) 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; $\tilde\kappa^2 \equiv 16\pi G \sim M_{\rm Pl}^{-2}$, 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; we quote the arXiv version of that reference, which is the version we were able to consult, and no discrepancy with the published version is known to us; 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. (2). 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 $\approx$3.3 at the LQG value $\gamma \approx 0.24$ and $\mathcal{O}(3)$–$\mathcal{O}(5)$ across $\gamma \lesssim \mathcal{O}(1)$ (the ratio is bounded below by $378/120 \approx 3.2$ for all real γ)], not a free normalization; and (iii) the explicit $\tilde{\kappa}^2 \sim M_{\rm Pl}^{-2}$ on every renormalized charge, i.e. the Planck suppression in the prefactor $\beta(\gamma)/M_{\rm Pl}$. 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 $\partial_\mu \vartheta_{\rm NY} J^{5\mu}$ 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. Because the closure below retains $\approx$60 (conservatively ≥ 58) orders of margin, it is insensitive to this residual $\mathcal{O}(1)$ ambiguity in $\beta(\gamma)$. The prefactor is $\mathcal{O}(\alpha_{\rm em}/4\pi)$ — a dimensionless loop factor — times a single inverse power of the Planck mass; the product carries mass dimension -1 and is not itself dimensionless. Explicitly, the prefactor $\beta(\gamma)/M_{\rm Pl}$ carries mass dimension -1 (the division by $M_{\rm Pl}$, not multiplication): with $\vartheta_{\rm NY}$ of dimension $+1$, $\partial_\mu\vartheta_{\rm NY}$ 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 $\partial_\mu\vartheta_{\rm NY} J^{5\mu}$ times the dimension-(-1) prefactor $\beta(\gamma)/M_{\rm Pl}$ written outside the integral in Eq. (2) — is a Lagrangian density of mass dimension $-1+2+3=+4$, so that the action, obtained by integrating it over d^4x, is dimensionless, as required. Once the dimension-$(+2)$ background value $\partial_\mu\vartheta_{\rm NY} \sim H_0^2 \sim 10^{-66}\,{\rm eV}^2$ at the present epoch is substituted (the cosmological Nieh–Yan pseudoscalar, of dimension $+1$ and Hubble-scale amplitude, evolving on the Hubble time), the one-loop rotation accumulated between recombination and today — over an elapsed time of order the Hubble time, $\Delta t \sim H_0^{-1}$ — assembles from the stated ingredients as the explicit intermediate estimate

$$\Delta\theta_{\text{one-loop}} \sim \frac{\alpha_{\rm em}}{4\pi}\,\frac{\partial_\mu\vartheta_{\rm NY}}{M_{\rm Pl}}\,\Delta t \sim \frac{\alpha_{\rm em}}{4\pi}\,\frac{H_0^2}{M_{\rm Pl}}\,H_0^{-1} = \frac{\alpha_{\rm em}}{4\pi}\,\frac{H_0}{M_{\rm Pl}}$$

(the anomaly-mediated photon chain of Appendix B supplies the $\alpha_{\rm em}/4\pi$; the one-loop factor $1/16\pi^2$ and the $\mathcal{O}(1)$ coefficient $\beta(\gamma)$ are conservatively dropped, their *net* effect being a further suppression of roughly 1.7 orders — the loop factor supplies -2.2 orders against $\beta(\gamma) \approx 3.3$'s $+0.5$, so the pair is dropped as a pair and not because each factor separately suppresses; this Tier-III accumulation ansatz is exactly the bookkeeping the evidentiary label of Table II records), and is stated as the dimensionless budget ratio

$$\frac{\Delta\theta_{\text{one-loop}}}{\beta_{\rm obs}\left[M_{\rm Pl}(\alpha/M)\right]} \sim \frac{\alpha_{\rm em}}{4\pi}\,\frac{H_0/M_{\rm Pl}}{M_{\rm Pl}\,(\alpha/M)\,\beta_{\rm obs}} \sim \frac{\alpha_{\rm em}}{4\pi}\,\frac{H_0}{M_{\rm Pl}}\,\frac{M}{\alpha\,M_{\rm Pl}\,\beta_{\rm obs}}\,, \tag{3}$$

whose left-hand side normalizes the one-loop amplitude *both* by the observed angle $\Delta\theta_{\rm obs} = \beta_{\rm obs}$ *and* by the dimensionless strength $M_{\rm Pl}(\alpha/M) \sim 10^{-2}$ of the R4-fitted coupling that reproduces that angle: the double normalization states the suppression relative to the fitted-coupling benchmark, the weaker (more conservative) of the two available normalizations, and both normalizations are evaluated explicitly below. Here $\alpha_{\rm em}/(4\pi) \approx 6\times10^{-4}$ (more precisely 5.8×10^{-4}, using the Thomson-limit fine-structure constant $\alpha_{\rm em} \simeq 1/137$; a running $\alpha_{\rm em}(\mu)$ evaluated at any higher scale differs by an $\mathcal{O}(1)$ factor immaterial at these margins), 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 $H_0/M_{\rm Pl} \sim 10^{-61}$, and the R4-fitted coupling $\alpha/M \sim 10^{-21}\,{\rm GeV}^{-1}$ (Sec. IV C) gives $M_{\rm Pl}\cdot(\alpha/M) \sim 10^{19}\,{\rm GeV}\cdot 10^{-21}\,{\rm GeV}^{-1} = 10^{-2}$. Plugging in $\beta_{\rm obs} = 0.342^\circ \approx 6\times10^{-3}\,{\rm rad}$ ($0.342\times\pi/180 = 5.97\times10^{-3}$), the dimensionless ratio Eq. (3) evaluates to $10^{-3}\cdot10^{-61}/(10^{-2}\cdot 6\times10^{-3}) \approx 10^{-60}$ (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 $1\,{\rm GeV} = 10^9\,{\rm eV}$ and is not a source of ambiguity). Contracting directly against the observed angle alone gives $(\alpha_{\rm em}/4\pi)(H_0/M_{\rm Pl})/\beta_{\rm obs} \approx 10^{-3}\cdot10^{-61}/(6\times10^{-3}) \approx 2\times10^{-62}$, *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 $\approx$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 $\mathcal{O}(1)$–$\mathcal{O}(10^{10})$ normalization ambiguity reopens the route. To reconcile the two registers in which this number is reported: Table II labels the birefringence leg ansatz-level and exploratory because the *precise size* of the margin rests on the scaling ansatz, whereas what the abstract's headline uses, and all this survey relies on, is only that the budget sits many orders below the observed amplitude — a conclusion the 48-order stress test shows is insensitive to the ansatz. 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 only a far looser upper bound and is not used in the closure. The conclusion is therefore robust to the bookkeeping choice: the one-loop Holst-sector parity-odd term sits far below not merely the WMAP+Planck (and ACT DR6) birefringence sensitivity but the observed central value itself, and cannot account for the observed angle. (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 (birefringence channel): amplitude-suppressed by* $M_{\rm Pl}^{-1}$ *and one-loop factor* $\alpha_{\rm em}/(4\pi)$.

Route 2's dark-energy leg, and the limits of its closure.—The budget just computed is a bound on a CMB rotation angle and says nothing about ρ_Λ. Its dark-energy leg must be stated separately, and honestly, because Eq. (2) is *not* a member of the dimension-four operator list of Sec. V. Three independent features of the operator put it outside that list's construction rule: it is a dimension-five operator, its dimension-(+5) integrand $\partial_\mu \vartheta_{\rm NY} J^{5\mu}$ carried by the dimension-(−1) coefficient $\beta(\gamma)/M_{\rm Pl}$ (whereas the list admits densities of operator dimension exactly four with *dimensionless* Wilson coefficients); the field $\vartheta_{\rm NY}(x)$ is not among the admitted building blocks (tetrad, torsionful curvature two-form, the algebraic torsion of Eq. (1), axial current J^5); and $\partial_\mu \vartheta_{\rm NY}$ carries a derivative beyond those internal to R and T, which the rule excludes. Its Hubble-scale background $\langle \partial_\mu \vartheta_{\rm NY} \rangle \sim H_0^2$ moreover makes $\vartheta_{\rm NY}$ a new light scale $\mu \ll M_{\rm Pl}$ — precisely the case App. A names as able to *evade* the single-scale NDA bound. We therefore do not claim the NDA bound covers it, and the Route-2 dark-energy leg splits:

(i) Constant coefficient (minimal ECH). In minimal ECH the Immirzi parameter, and with it the Nieh–Yan coefficient, is a constant fixed by the LQG area spectrum, with no landscape for it to vary over (B7). For constant $\vartheta_{\rm NY}$, $\partial_\mu \vartheta_{\rm NY} = 0$ and Eq. (2) vanishes identically; the surviving Holst/Nieh–Yan content is the pair O1/O2 of the Sec. V list. O2 is an exact total derivative, contributing zero to the equations of motion and zero to the vacuum energy; on the on-shell branch O1 equals −O2 plus $\frac{1}{2}\mathcal{O}_4^{[4]}$ [Eq. (12)], and $\mathcal{O}_4^{[4]}$ is the Fierz-closed $(J^5 \cdot J^5)$ contact term of Eq. (13), bounded at the same $M_{\rm Pl}^{-2}$ power as O5 by the projection lemma of Appendix C and the single-scale ceiling of Appendix A. Minimal Route 2 therefore sources no dark energy beyond that Planck-suppressed contact term. The operator-level input to that conclusion — that O2 is an exact total derivative and that O1's remainder is Fierz-closed and $M_{\rm Pl}^{-2}$-suppressed — is rigorous (Tier-I) about the list. The conclusion itself needs one step more: that O1/O2 exhaust the rule-admitted Holst/Nieh–Yan content. That is the spanning assertion, which this survey states from its construction rules and does *not* prove by exhaustive symbolic enumeration, so the conclusion inherits it and the leg is recorded at Tier-II in Table II. The catalog's only Tier-I leg remains the perturbation-transparency theorem of Appendix D.

(ii) Dynamical coefficient (non-minimal). The only way the Nieh–Yan sector can source dynamics is a non-constant $\vartheta_{\rm NY}$ carrying its own kinetic term or potential — i.e. a dynamical-Immirzi-type completion, which Sec. V, App. A 1 and Sec. VI all place outside the minimal scope. A $\vartheta_{\rm NY}$ with an $\sim H_0$ mass/potential tuned to yield ρ_Λ is Route 4 in gravitational costume: it re-imports exactly the cosmological-constant fine-tuning that closes R4 at the naturalness / explanatory-deficit level (Sec. IV C), and B7 denies minimal ECH the varying γ such a field would require. On the scalar-matter branch the induced parity channel is additionally the channel shown inert by B14 (App. D).

We state the resulting strength plainly. Route 2's dark-energy leg is closed at Tier-II in both cases: in case (i) by an operator-level argument whose rigorous ingredient (O2 is an exact total derivative and O1's on-shell remainder is a Fierz-closed $M_{\rm Pl}^{-2}$-suppressed contact term) is Tier-I but whose conclusion inherits the unproved spanning assertion, and in case (ii) by the same naturalness objection that closes R4 — and in neither case by an amplitude bound or by the single-scale NDA argument. No dark-energy amplitude is computed for Route 2 anywhere in this survey, and none is claimed. The evidentiary status of the estimate, stated once: Eq. (2) is an *illustrative upper-bound amplitude budget* constructed as the natural EFT operator at the $M_{\rm Pl}^{-1}\alpha_{\rm em}/(4\pi)$ scale, *not* a result extracted from Mercuri [8] or Date–Kaul–Sengupta [9], 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 [9]; 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

$$\frac{d\gamma}{d\ln\mu} = \frac{1}{12\pi^2}\,(N_F^L - N_F^R)\,\gamma + \mathcal{O}(\gamma^2), \tag{4}$$

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. (4) only as an upper-bound EFT ansatz for the Route-3 amplitude budget and do not claim it is taken verbatim from [9]. The actual fermion-induced perturbative running of the Immirzi parameter is computed by Benedetti & Speziale [10], 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. (4) 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; we quote it from their proceedings summary [3] (whose Eq. 7 is the equation numbering used here), companion to the full one-loop analysis of Ref. [10]:

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

whose only fixed point (of this Benedetti–Speziale flow; the fixed-point-free Riccati system above is Shapiro–Teixeira's Route-2 λ_4–γ 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 $\mu^2\tilde{\kappa}^2 \simeq (\mu/M_{\rm Pl})^2$ 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 $(\mu_{\rm UV}/M_{\rm Pl})^2$. Integrating Eq. (5) numerically from a GUT-scale ultraviolet boundary $\mu_{\rm UV} \sim 10^{16}\,$GeV down to $\mu_{\rm IR} \sim 1\,$GeV (with γ near the LQG value $\gamma \approx 0.24$) gives, in agreement with a frozen-coefficient analytic estimate, $|\Delta\gamma/\gamma| \approx 1.4\times10^{-6}$ (evaluated with the full $M_{\rm Pl} \simeq 1.22\times10^{19}\,$GeV, per the Sec. II convention for order-of-magnitude numerics). This is a genuine integrated running, *not* an ansatz bound: the explicit $\mu^2\tilde{\kappa}^2 \simeq (\mu/M_{\rm Pl})^2$ prefactor makes the flow power-suppressed, so $|\Delta\gamma/\gamma| \sim (\mu_{\rm UV}/M_{\rm Pl})^2$ 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 $|\Delta\gamma/\gamma| \to \mathcal{O}(1)$, 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 mass-dimension scaling relation Eq. (6) displayed below, the derived torsion/Immirzi contribution to ρ_Λ sits $\sim$ 61–67 orders of magnitude below the observed dark-energy density $\rho_{\Lambda,\rm obs} \approx (2.25\,\text{meV})^4$ — 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. (4), $\Delta\gamma/\gamma \sim 0.3$, 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. (4), which is linear in γ, gives $\Delta\ln\gamma \approx (N_F^L - N_F^R)\,\ln(\mu_{\rm GUT}/\mu_{\rm IR})/(12\pi^2)$; with a net chiral count $(N_F^L - N_F^R) = \mathcal{O}(1)$ and a GUT-to-IR lever arm $\ln(\mu_{\rm GUT}/\mu_{\rm IR}) \approx$ 30–37 ($\mu_{\rm GUT} \sim 10^{16}\,$GeV throughout: the upper edge from running down to $\mu_{\rm IR} \sim 1\,$GeV, $\ln 10^{16} \approx 36.8$; the lower edge from cutting the flow at $\mu_{\rm IR} \sim 1\,$TeV, i.e. crediting only running above collider-probed scales, $\ln 10^{13} \approx 30$), this is numerically $\Delta\ln\gamma \approx$ 0.25–0.31 (a few $\times10^{-1}$; e.g. $32/(12\pi^2) \approx 0.27$; we identify $\Delta\gamma/\gamma \simeq \Delta\ln\gamma$ at this order — exponentiating would give 0.29–0.36, a distinction immaterial at the $\gtrsim$ 60-order margins below). 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, and *not* claimed to be structurally consistent with the Benedetti–Speziale flow. The two flows are structurally different: Eq. (4) is purely logarithmic with its only fixed point at $\gamma = 0$, whereas Eq. (5) is power-suppressed by $\mu^2\tilde{\kappa}^2$ with its sole fixed point at $\gamma^2 = 1$ and sign set by $|\gamma| \gtrless 1$. What is defensible, and all that is used, is the *numerical* inequality $0.3 \gg 1.4\times10^{-6}$: the chiral-count value numerically bounds the derived integrated running from above — 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 $M_{\rm Pl}^{-1}$ in the prefactor in any cosmologically relevant scalar-curvature regime. Writing that scaling out explicitly, the induced parity-odd contribution to the late-time vacuum energy is

$$\frac{\rho_{\rm R3}}{\rho_{\Lambda,\rm obs}} \sim \left(\frac{\Delta\gamma}{\gamma}\right)\left(\frac{H_0}{M_{\rm Pl}}\right), \qquad (6)$$

where the reference budget in the denominator is the observed dark-energy density itself, $\rho_{\Lambda,\rm obs} \approx (2.25\,\text{meV})^4$: the single inverse Planck power fixed by the dimension count is saturated, in a late-time cosmological background, by the only available infrared scale, the present Hubble rate H_0, and $\Delta\gamma/\gamma$ is the dimensionless amplitude the running supplies. Eq. (6) is a Tier-III mass-dimension scaling relation, not a derived stress-tensor matching — we do not supply the cosmological stress tensor that a genuine mapping would require, and none is claimed — but it is what the "61–67 orders" statement means, and it is stated here rather than left implicit. Numerically, with $H_0/M_{\rm Pl} = 1.18\times10^{-61}$: the conservative chiral-count input $\Delta\gamma/\gamma \sim 0.3$ gives $\rho_{\rm R3}/\rho_{\Lambda,\rm obs} \sim 3\times10^{-62}$ ($\sim$61 orders), and the derived integrated value $|\Delta\gamma/\gamma| \approx 1.4\times10^{-6}$ gives $\sim 1.7\times10^{-67}$ ($\sim$67 orders) — the two endpoints quoted in the abstract. Both close this route by many orders of magnitude. Note that H_0, not a bounce-epoch H, is the correct insertion throughout: at the bounce $H/M_{\rm Pl} = \mathcal{O}(1)$, and the claim being made is about the *late-time* dark-energy channel. *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. (5) 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 $\tilde{\kappa}^2 \sim M_{\rm Pl}^{-2}$ 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 *birefringence* amplitude is bounded by the explicit budget of Eq. (3) regardless of that $\mathcal{O}(1)$ normalization. (The single-scale NDA bound is *not* claimed to cover Eq. (2), whose H_0-scale pseudoscalar background is exactly the light scale App. A names as able to evade it; see Sec. IV A.) The kept ansatz value Eq. (4) ($\Delta\gamma/\gamma \sim 0.3$) 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 one-loop parity-odd 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. (4) 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, $\mathcal{L}_{\rm CS} \supset -\frac{1}{4}(\alpha/M)\,\phi\,\widetilde{F}_{\mu\nu}F^{\mu\nu}$, producing a rotation of the CMB polarization plane (cosmic birefringence). Fitting this coupling to the observed WMAP+Planck birefringence signal $\beta_{\rm obs} = 0.342^\circ \pm 0.094^\circ$ (WMAP+Planck [11]; the first Planck-2018 extraction reported $0.35^\circ \pm 0.14^\circ$ [12]) (comparable to the independent ACT DR6 measurement $\beta = 0.215^\circ \pm 0.074^\circ$ [13]) gives the anchor value $\alpha/M \sim 10^{-21}\,{\rm GeV}^{-1}$ used above as a numerical input to the Route-2 amplitude estimate. Both measurements are statistical indications rather than established detections — $\approx 3.6\sigma$ for WMAP+Planck and $\approx 2.9\sigma$ for ACT DR6, significances obtained from different datasets and distinct null procedures and therefore not directly comparable as statistical weights — and the Route-2 suppression conclusion is insensitive to that status: a smaller or vanishing true birefringence signal only *widens* the margin of Eq. (3). The algebraic origin of that scale, carried from the companion's derivation [1], is the standard small-rotation result $\beta = (\alpha/2M)\,\Delta\phi_{\rm rec\to today}$ (the rotation angle is half the coupling times the coherent field excursion), with excursion $\Delta\phi \sim \sqrt{2\rho_\theta}/m_\theta \sim M_{\rm Pl}$ for a frozen ($m_\theta \lesssim H_0$) spectator carrying $\rho_\theta \sim \rho_\Lambda$; matching $\beta_{\rm obs} \approx 6\times10^{-3}$ rad then gives $\alpha/M \sim 2\beta_{\rm obs}/M_{\rm Pl} \sim 10^{-21}\,{\rm GeV}^{-1}$. R4 does not close by amplitude suppression: the same coupling that reproduces $\beta_{\rm obs}$ 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. The two steps a referee needs are given above and are checkable here: the coupling estimate $\alpha/M \sim 2\beta_{\rm obs}/M_{\rm Pl}$ and the observation that neither $m_\theta \sim H_0$ nor that value of α/M follows from the minimal ECH field content. What is *not* reproduced here is the full spectator-ALP derivation and the companion's consistency check against a birefringence benchmark [1, 14]; the R4 leg of this survey's accounting is scoped accordingly, and rests on those external, not-yet-peer-reviewed derivations for anything beyond the naturalness statement made here.

Within the four-channel enumeration of Sec. II, Routes R1–R4 cover the four parity-odd or dark-energy channels of this framework. Two operators occasionally raised as omissions from this four-channel enumeration — the Jackiw–Pi gravitational Chern–Simons term $R\wedge\widetilde{R}$ and the parity-odd four-fermion partner of R1 carrying the $\gamma_{\rm BI}/(\gamma_{\rm BI}^2{+}1)\cdot 8\pi G$ 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 $M_{\rm Pl}^{-2}$ and parity-even. R2 (one-loop graviton corrections) is amplitude-suppressed by $M_{\rm Pl}^{-1}$ and the one-loop factor $\alpha_{\rm em}/(4\pi)$ *in its birefringence channel*; its dark-energy leg is closed operator-level for constant Nieh–Yan coefficient and only at the R4-class naturalness level for a dynamical one (Sec. IV A). 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 $\beta_{\rm obs}$ 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* in the channel each is bounded against (the predicted amplitude is suppressed by explicit $M_{\rm Pl}$ powers and loop factors, placing it many orders below any observable threshold regardless of normalisation assumptions); the one exception, stated rather than absorbed, is R2's dark-energy leg on the dynamical-$\vartheta_{\rm NY}$ branch, which closes in the naturalness mode below rather than the amplitude mode. 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 $\beta_{\rm obs}$, with the dark-energy density supplied by an unrelated sector (e.g. a quintessence field [15], a quintom-class two-field scenario [16], 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-LIST 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 space is considered. This section answers that question: we exhibit a finite *spanning list* 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 list, not to one representative. We stress at the outset that the list is a generating set, not a linearly independent basis: it carries two exact relations, quantified in Eq. (12) below, and is retained in redundant form because each entry is a separately *recognizable* invariant of the literature. Bounding every member of a spanning list bounds every density it spans, which is all the no-go requires; linear independence is neither claimed nor needed.

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 [8] — consider the parity-odd term

$$S_{\rm eff} = \frac{\alpha}{M}\int e_I \wedge e_J \wedge \mathcal{F}^{IJ}[K, \mathring{R}], \tag{7}$$

where $M = M_{\text{area-gap}} \sim M_{\rm Pl}/\sqrt{\gamma}$ is the LQG area-gap mass scale (from the LQG area-gap $\Delta \propto \gamma\,\ell_P^2$, the inverse-length / mass scale is $M_\Delta \sim M_{\rm Pl}/\sqrt{\gamma}$ up to numerical constants), α is a dimensionless coupling, and $\mathcal{F}^{IJ}[K, \mathring{R}]$ denotes the curvature two-form of the full (torsionful) connection, written as a functional of the contorsion K and the Levi-Civita curvature $\mathring{R}$; the component form Eq. (8) displays its leading contribution.

In components, the leading contribution reduces to

$$S_{\rm eff} = \int d^4x\,\sqrt{-g}\,\frac{\alpha}{M}\,\varepsilon^{\mu\nu\rho\sigma}\,e^I_\mu\,e^J_\nu\,\mathcal{F}_{IJ\rho\sigma}, \tag{8}$$

which has naive mass dimension $[\mathcal{L}_{\rm odd}] = +1$—three units short of the required $+4$ (with $[\alpha] = 0$, $[M] = +1$, $[e^I_\mu] = 0$, 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.

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

$$S^{(4)}_{\rm eff} = \int d^4x\,\sqrt{-g}\sum_n c_n\,\mathcal{O}^{[4]}_n, \qquad \{\mathcal{O}^{[4]}_n\} = \{\text{O1–O6}\}, \tag{9}$$

where each $\mathcal{O}^{[4]}_n$ is a genuine dimension-4 Lagrangian density and each Wilson coefficient c_n is a dimensionless rational. Throughout this paper "genuine dimension-4"

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 $M_{\rm Pl}^{-2}$ amplitude suppression is a Tier-III order-of-magnitude estimate from a standard Cartan torsion-elimination derivation: the companion's coefficient-one benchmark $\kappa n_\psi^2/\rho_\Lambda \simeq 3.6 \times 10^{-69}\,(n_\psi/100\,{\rm cm}^{-3})^2$, i.e. $\approx$68 orders below ρ_Λ even at a deliberately elevated ISM-like normalization [1] (arithmetic carried self-contained in App. E 2, Eq. (E5)).
R2 (one-loop graviton corrections)	One-loop promotion of $\gamma_{\rm BI}$ reaching the observed birefringence amplitude (and, separately, ρ_Λ)	**(III) Ansatz-level** for the *birefringence* leg: amplitude budget suppressed by $H_0/M_{\rm Pl}$ and $\alpha_{\rm em}/(4\pi)$ under an explicitly-labeled scaling ansatz; exploratory in its precise size, not in its sign (Sec. IV A). *Dark-energy* leg: **(II)** in both branches. Constant Nieh–Yan coefficient (B7 fixes γ): O2 is an exact total derivative and O1's remainder is Fierz-closed and $M_{\rm Pl}^{-2}$-suppressed, a Tier-I fact about the list, but the step to "R2 sources no dark energy" inherits the unproved spanning assertion (Sec. IV A). Dynamical $\vartheta_{\rm NY}$: R4-class naturalness. No dark-energy amplitude is computed for R2.
R3 (Immirzi running)	Quantum running of $\gamma_{\rm BI}$ generating a dark-energy-relevant shift	**(II)+(III)**: mass-dimension lock is structural; the chiral-count ansatz bound (motivated by Date–Kaul–Sengupta [9]) 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 *deriving* the single coupling that yields both $\beta_{\rm obs}$ and ρ_Λ	**(II) Structural naturalness objection**, *not* an amplitude exclusion: a free-α/M ALP fit reproduces $\beta_{\rm obs}$, 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).

is reserved for the promoted operators $\mathcal{O}_n^{[4]}$, *all six* of which carry mass dimension exactly +4 by construction;

it never refers to the subset of *bare* invariants that happen to sit at dimension +4 before promotion (O3 and O5), nor to the subset that reduces onto the Fierz basis. The sum runs over the six named invariants, which span — but do not independently span — the rule-admitted space. This off-shell dimension-4 operator space is the *primary physical foundation* of the no-go: the single-scale NDA ceiling, the Fierz closure, and the Bianchi-vanishing results below are all derived properties of it, and none depends on the on-shell shorthand Eq. (8) for its logical force. *Construction rule.*—The admitted building blocks are the tetrad e^I_μ (with the invariant tensors ε and η), the curvature two-form of the torsionful connection, the algebraically-fixed torsion of Eq. (1), 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 $M^2_{\rm Pl}$ promotions of Eq. (10). Under this rule the admitted densities are spanned by the six-member list of Eq. (10). That the list *spans* the rule-admitted space 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 scripts verify the two reduction identities used below and the two exact relations of Eq. (12), 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 contains, already *within* the minimal field content, the gravitational contribution to the chiral anomaly, $\nabla_\mu J^{5\mu} \supset c_{\rm grav}\,\varepsilon^{\mu\nu\rho\sigma}R_{\mu\nu\alpha\beta}R_{\rho\sigma}{}^{\alpha\beta}$ with a known pure-number coefficient (Kimura–Delbourgo–Salam [17, 18]) — but that density is exactly the Pontryagin invariant O3, already a member of the spanning list below and disposed of as an exact total derivative; the $F\widetilde{F}$ anomaly content arises 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 $[S] = [J^5] = +3$, $[\kappa] = [M^{-2}_{\rm Pl}] = -2$), the operator *definitions* include the explicit compensating $M_{\rm Pl}$ 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 spanning list.*—Eq. (8) with its dimension-+1 integrand is a phenomenological on-shell representative of the dimension-4 spanning list in which the coefficient α/M (mass dimension −1) stands in for the listed operators' $M^2_{\rm Pl}$ 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: the apparent dimension-+1 deficit is entirely that α/M-for-$M^2_{\rm Pl}$ coefficient substitution (Appendix A), not any ill-defined operator. The Case II on-shell curvature dressing (inserting $R \sim M^2_{\rm Pl}$ to close the dimension count) is an *illustrative heuristic*, subordinate to the operator list 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)

$$\begin{aligned}
\mathcal{O}^{[4]}_1 &= M^2_{\rm Pl}\,\varepsilon\, e^I e^J R_{IJ}, & \mathcal{O}^{[4]}_2 &= M^2_{\rm Pl}\,{\rm NY},\\
\mathcal{O}^{[4]}_3 &= R^{IJ}\wedge R_{IJ}, & \mathcal{O}^{[4]}_4 &= M^2_{\rm Pl}\,\varepsilon^{\mu\nu\rho\sigma}T^I{}_{\mu\nu}T_{I\rho\sigma},\\
\mathcal{O}^{[4]}_5 &= \varepsilon\, T\, e\, J^5, & \mathcal{O}^{[4]}_6 &= M^2_{\rm Pl}\,\varepsilon^{\mu\nu\rho\sigma}R_{\mu\nu\rho\sigma},
\end{aligned} \tag{10}$$

where $\varepsilon^{\mu\nu\rho\sigma}$ throughout is the spacetime Levi-Civita *tensor* with $\varepsilon^{0123} = +1$, equivalently $\varepsilon_{0123} = -1$ (mostly-plus signature, Sec. II) and R is the curvature of the *full torsionful* connection per the construction rule. The two schematic entries are the unique admitted full contractions of their factors against one ε, namely $\mathcal{O}^{[4]}_1 = M^2_{\rm Pl}\varepsilon^{\mu\nu\rho\sigma}e^I_\mu e^J_\nu R_{IJ\rho\sigma}$ (an internal ε_{IJKL} is inadmissible here — I, J are already contracted between the tetrads and R_{IJ}, and that variant would be the parity-*even* Palatini term) and $\mathcal{O}^{[4]}_5 = \varepsilon^{\mu\nu\rho\sigma}T^I{}_{\mu\nu}e_{I\rho}J^5_\sigma$, the unique zero-derivative contraction, whose dimensions $+1+0+3 = +4$ match the count below. The Nieh–Yan entry is fixed in the *density* normalization

$${\rm NY} \;\equiv\; \partial_\mu\big(\varepsilon^{\mu\nu\rho\sigma}e_{I\nu}T^I{}_{\rho\sigma}\big)\,. \tag{11}$$

Stating Eq. (11) explicitly matters: the form-versus-density normalization of the Nieh–Yan entry is what fixes the rational coefficients in Eq. (12) below, and a schematic normalization misstates them by a factor of two. The conversion is one line. Writing $\mathcal{A} \equiv \varepsilon^{\mu\nu\rho\sigma}e^I_\mu e^J_\nu R_{IJ\rho\sigma}$ and $\mathcal{B} \equiv \varepsilon^{\mu\nu\rho\sigma}T^I{}_{\mu\nu}T_{I\rho\sigma}$ for the bare invariants underlying O1 and O4, the strict form-to-density conversion of the Nieh–Yan identity gives $[d(e_I\wedge T^I)]_{\rm dens} = \frac{1}{4}\mathcal{B} - \frac{1}{2}\mathcal{A}$, whereas Eq. (11) defines ${\rm NY} = \frac{1}{2}\mathcal{B} - \mathcal{A}$, twice that. The relative factor between the $T\wedge T$ and $e\wedge e\wedge R$ terms is wedge combinatorics, not a rescaling of NY alone, which is why the schematic 1 : 1 : 1 reading of the identity does not reproduce Eq. (12). Here the torsion two-form T^I carries a single internal index (components $T^I{}_{\mu\nu}$; the component tensor T^{abc} of Eq. (1) 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 $d(e_I \wedge T^I) = T_I \wedge T^I - e_I \wedge e_J \wedge R^{IJ}$ [19]. The list 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 $M^2_{\rm Pl}$ 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 $M_{\rm Pl}^2$ prefactors explicitly so that every $c_n\mathcal{O}_n^{[4]}$ is a bona-fide dimension-4 density.

The list is a spanning list, not a basis.—The six named invariants are not linearly independent. Expanding all six over a common basis of independent jet monomials, built from the Cartan structure equations on an algebraically independent 2-jet of $(e, \partial e, \partial\partial e, \omega, \partial\omega, J^5)$ and reduced in exact rational arithmetic, gives a coefficient matrix of *rank four* with a two-dimensional null space, spanned by

$$\mathcal{O}_1^{[4]} - \mathcal{O}_6^{[4]} = 0, \qquad 2\,\mathcal{O}_1^{[4]} + 2\,\mathcal{O}_2^{[4]} - \mathcal{O}_4^{[4]} = 0, \tag{12}$$

the second equivalently $\mathcal{O}_1^{[4]} = \frac{1}{2}\mathcal{O}_4^{[4]} - \mathcal{O}_2^{[4]}$. Both hold identically off shell and on shell. The first is the tetrad conversion $e^I_\mu e^J_\nu R_{IJ\rho\sigma} = R_{\mu\nu\rho\sigma}$ — the same conversion this paper performs for Eq. (8) — so O1 and O6 are literally the same density; the second is the Nieh–Yan identity in the normalization Eq. (11), whose rational coefficients are normalization-dependent even though the existence of the relation is not. Because O2 and O3 are exact forms, the rank drops further, to *two* modulo total derivatives. Independent subsets of maximal rank include $\{$O2, O3, O4, O5$\}$ and $\{$O1, O3, O4, O5$\}$. The redundancy is deliberate and is retained so that each entry remains a separately recognizable invariant of the literature (Holst dual, Nieh–Yan, Pontryagin, torsion-square, axial–torsion, single-curvature); the no-go statement is a statement about the space these six *span* in Eq. (9), and a bound holding on every member of a spanning list holds on every density in the span. These relations, the rank, and the null space were obtained by independent symbolic computation (the independent operator-basis adjudication script, listed with its repository location in the Data and Code Availability statement), not by rearranging the identities quoted above.

The physics conclusion is unchanged: the algebraic Cartan–Holst constraint [Eq. (1)], the first (algebraic) Bianchi identity, and Chern–Weil exactness collapse this set to (i) topological total derivatives (O2 and O3), (ii) the $M_{\rm Pl}^{-2}$-suppressed Fierz-closed four-fermion sector (O5, O4, and the non-total-derivative remainder of O1 = O6 on the on-shell branch), or (iii) densities that vanish identically (O1 and O6 on the torsion-free branch, by the algebraic Bianchi identity). The ε-contracted torsion-square is supported only by the non-axial torsion irreps, and the on-shell torsion of Eq. (1) supplies one: O4 is nonzero on shell and lands, like O5, in class (ii).

Main-text closure argument.—The collapse is closed by the exhibited spanning list plus two tensor identities; the identities are verified symbolically in the released script `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. That branch qualifier is essential: the identity used is the *torsion-free* first Bianchi identity, which a nonzero on-shell torsion violates. On the on-shell branch O1 and O6 are not pointwise zero, and Eq. (12) gives $\mathcal{O}_1^{[4]} = \mathcal{O}_6^{[4]} = -\mathcal{O}_2^{[4]} + \frac{1}{2}\mathcal{O}_4^{[4]}$ exactly: minus the Nieh–Yan total derivative, which contributes zero to the equations of motion and zero to the vacuum energy, *plus* half the ε-contracted torsion-square, which is the Fierz-closed $M_{\rm Pl}^{-2}$-suppressed contact term of Eq. (13). O1 and O6 are therefore not exact total derivatives on shell; their remainder sits in the same disposal class as O5. (b) The axial–torsion density O5 reduces under Eq. (1), using $S^{abc} = \frac{1}{4}\varepsilon^{abcd}J^5_d$ and the Lorentzian $\varepsilon_{abcd}\varepsilon^{abce} = -3!\,\delta^e_d$ (consistent with the companion paper's torsion-elimination normalization [1]), to the four-fermion contact operator $-3\kappa\,[\gamma^2/(1+\gamma^2)]\,(J^5\cdot J^5)$ — itself parity-even, since a Lorentz-scalar product of two axial currents is P-even (B8, Sec. III); the parity-odd label belongs to the pre-reduction ε-contracted densities — 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 $M_{\rm Pl}^{-2}$ power. Only the axial irrep contributes: the trace-vector part of Eq. (1) drops out of O5 entirely. The ε-contracted torsion-square O4 joins O5 rather than vanishing. $T_I\wedge T^I$ is supported only by the non-axial torsion irreps — it vanishes on a pure axial and on a pure trace-vector torsion alike, and is carried entirely by the axial × trace-vector cross term — and the on-shell torsion of Eq. (1) populates exactly that channel, giving

$$\mathcal{O}_4^{[4]} = -24\,M_{\rm Pl}^2\,\alpha\beta\,(J^5\cdot J^5) = -\frac{3\kappa\,\gamma^3}{(1+\gamma^2)^2}\,(J^5\cdot J^5), \tag{13}$$

whose bare invariant is $-192\pi^2G^2[\gamma^3/(1+\gamma^2)^2](J^5\cdot J^5)$. This is the same Fierz-closed $(J^5\cdot J^5)$ structure as O5 at the same $M_{\rm Pl}^{-2}$ power, with $\mathcal{O}_4^{[4]}/\mathcal{O}_5^{[4]} = \gamma/(1+\gamma^2) \simeq 0.22$ at $\gamma = 0.2375$; it vanishes only in the $\gamma\to\infty$ Einstein–Cartan limit, where the trace-vector irrep switches off. The same on-shell torsion also renders the ε-free torsion–current density

$$T^a{}_{ab}\,J^{5b} = 3\beta\,(J^5\cdot J^5), \qquad \beta = \frac{\kappa\gamma}{4(1+\gamma^2)}, \tag{14}$$

nonzero (the axial irrep drops out of the trace, $\varepsilon^{ab}{}_{cd}$ being antisymmetric in the contracted pair, so only the trace-vector irrep survives, with $\eta^{ab}(\eta_{ba}J^5_c - \eta_{bc}J^5_a) = 3J^5_c$). Before torsion elimination this density is parity-odd — a polar torsion trace vector contracted with the axial current — yet it carries no ε tensor, so it is not reached by the ε-contraction construction rule. Its value is the same Fierz-closed $(J^5\cdot J^5)$ structure at the same $M_{\rm Pl}^{-2}$ power as O4 and O5, so it joins that bounded class and enlarges neither the closure argument's scope nor its conclusion; the enumeration claim below is therefore stated for the ε-contracted densities the construction rule generates, with this ε-free companion carried explicitly alongside them. (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 identically vanishing — and the Fierz-basis-reducible class, now carrying O4 and the O1 = O6 remainder alongside O5, is bounded uniformly at $M_{\rm Pl}^{-2}$ — so none produces a $(\mathrm{meV})^4$ vacuum energy without a new light scale $\mu \ll M_{\rm Pl}$ 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 A 1–C (Table III). On the on-shell branch (algebraic torsion eliminated; O2/O3 total derivatives; O4, O5 and the O1 = O6 remainder Fierz-closed and $M_{\rm Pl}^{-2}$-suppressed by Eq. (12) and Eq. (13)), this closed dimension-4 set is represented by the single dimension-+1 shorthand of Eq. (8), whose coefficient α/M (dimension -1) replaces the listed operators' $M_{\rm Pl}^2$ 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 $\rho_\Lambda = \Xi\, M_{\rm Pl}^4$ is therefore governed by a *single-scale NDA dimensional no-go* that holds identically for the formal dimension-4 operator list and its dimension-+1 shorthand (Appendix A 1 verifies the closure survives at genuine dimension 4 without the on-shell dressing): the three-unit deficit forces the missing powers to $M_{\rm Pl}^3$ (the only scale in minimal ECH), so the natural density is $\sim M_{\rm Pl}^4$, never $(\mathrm{meV})^4$ (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. Fourteen 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. Route 2's amplitude budget is a birefringence bound and is not a dark-energy bound: no dark-energy amplitude is computed for Route 2 in this survey, and its dark-energy leg closes operator-level only for a constant Nieh–Yan coefficient, and at the R4-class naturalness level otherwise (Sec. IV A). The amplitude closures are robust to $\mathcal{O}(1)$–$\mathcal{O}(10^{10})$ rescalings of their ansatz-level normalizations (Sec. IV A displays the stress test; applied to Route 3's pessimistic lower endpoint it leaves $\gtrsim 51$ orders). The operator-list argument (Sec. V) extends the governing dimensional bound from one representative operator to the exhibited six-member local dimension-four parity-odd operator spanning list {O1–O6} admitted by the minimal-coupling field content, with no escape operator identified within its span; the list is rank-four and deliberately redundant [Eq. (12)], and that it *spans* the rule-admitted space 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-list 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 the tensor (16) torsion irrep and the other torsion couplings that only a relaxation of minimal coupling admits — any of these constitutes a genuine escape from the minimal-coupling scope, not a loophole within it. The trace-vector (4) irrep is *not* in that excluded set: it is generated by the Holst term under strictly minimal coupling [Eq. (1)] and is carried throughout, entering the operator list through $\mathcal{O}_4^{[4]}$ [Eq. (13)], through the O1 = O6 remainder, and through the ε-free density of Eq. (14). A non-minimal completion introducing a new light scale $\mu \ll M_{\rm Pl}$, 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 fourteen 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, and only within the zero-spin branch its hypotheses cover: it excludes fermion sources, quantum loops and anomalies, non-minimal matter, a dynamical Immirzi field, and propagating torsion, so it bears on R2–R4 as a statement that their classical zero-spin baseline is inert and does *not* constrain the fermionic content of R1 (that is B8's independent role). 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, $\beta_{\rm obs}$) 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 14 distinct mechanism-class constraints — one per catalog entry — 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-list argument extends the governing single-scale dimensional bound from one representative parity-odd operator to the exhibited six-member local dimension-four operator spanning list {O1–O6} admitted by the minimal-coupling field content under stated construction rules: every member of that list is topological, Fierz-basis-reducible, or identically vanishing, with no identified escape operator within its span and within the stated minimal scope (the list is rank-four, hence a deliberately redundant generating set rather than a linearly independent basis; 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 under the stated minimal-coupling assumptions — R1 and R3 by amplitude suppression against ρ_Λ, R2 by amplitude suppression in its birefringence channel together with an operator-level disposal of its dark-energy leg for a constant Nieh–Yan coefficient, and R4 (and R2's dark-energy leg on the non-minimal dynamical-pseudoscalar branch) by naturalness / explanatory deficit rather than by an amplitude exclusion.

Data and Code Availability

The symbolic verification scripts supporting the operator-list 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
research/theory_audit/operator_basis_adjudication_2026_08_07.py
research/theory_audit/operator_basis_adjudication_2026_08_07.md
research/theory_audit/ech_torsion_onshell_2026_08_08.py
research/theory_audit/ech_torsion_onshell_2026_08_08.md
```

The first script verifies the two reduction identities of Appendix A 1 (Checks A and D); the second verifies the companion paper's Fierz convention (see the convention note in Appendix C); the third independently verifies 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. The fifth script is an independent symbolic adjudication of the dimension-four operator list itself: it rederives O1–O6 from the Cartan structure equations on an algebraically independent 2-jet, computes the exact rational rank and null space that give Eq. (12), certifies O1 = O6 by an independent affine-connection route, and evaluates every member on an explicitly curved on-shell Einstein–Cartan ($\gamma \to \infty$) configuration; the sixth file is its human-readable report. The seventh script solves the Holst-modified connection equation itself at finite γ — varying the first-order Einstein–Cartan–Holst action with respect to all 24 independent contorsion components with no irrep ansatz imposed, and cross-checking against an independent differential-form route — and returns the on-shell torsion irrep content of Eq. (1), the closed forms for O4 and O5, and the values of O1–O6 on six explicitly curved on-shell Einstein–Cartan–Holst configurations at $\gamma \in \{19/80, 1, 3\}$; the eighth file is its human-readable report. None of the scripts performs the enumeration establishing that the list spans the rule-admitted space, which rests on the stated construction rule. The first five files are frozen at immutable repository commit `1130b7c5e3d2`, whose copies are byte-identical to the current repository head; the operator-list adjudication report and the two on-shell torsion files postdate that commit and are available at the repository head. 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 (CC-BY-4.0) [1], whereas the third through eighth files listed above (the independent Fierz-adjudication, operator-list adjudication and on-shell torsion modules and their human-readable reports) are not part of that deposit and are available in the repository. An updated archival deposit containing this survey's own verification scripts is planned prior to publication, so that every computational artifact cited here is backed by a frozen-release DOI.

ACKNOWLEDGMENTS

This survey builds on the published spin-torsion cosmology of Nikodem Popławski [20], on the published derivations of Simone Mercuri [8] and of Laurent Freidel, Djordje Minic, and Tatsu Takeuchi [7] 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. 8) has off-shell mass dimension +1, not the +4 required for a local Lagrangian density:

$$[\alpha/M] = -1, \quad [\varepsilon^{\mu\nu\rho\sigma} e^I_\mu e^J_\nu \mathcal{F}_{IJ\rho\sigma}] = +2 \implies [\mathcal{L}_{\rm odd}] = +1. \tag{A1}$$

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 [21–23]. In minimal ECH the *only* available scale is $\Lambda \sim M_{\rm Pl}$: there is no intermediate threshold and no assumed cancellation. The three missing mass powers are therefore *forced* to be $M^3_{\rm Pl}$ — one cannot conjure a light scale from a theory that has none — and there are exactly two dimensionally-consistent ways to supply them.

Case I — coefficient dressing (local NDA reading).—The relevant operator here is the dimension-(+1) static piece inside $\varepsilon\, e\, e\, \mathcal{F}$, not the dimension-(+2) contraction itself; its natural coefficient is $c \sim M^{4-1}_{\rm Pl} = M^3_{\rm Pl}$, 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 $\rho \sim M^3_{\rm Pl}\langle T\rangle$. A coherent cosmological torsion is forbidden today by parity and isotropy ($\langle T\rangle \to 0$, hence $\rho \to 0$), while at the bounce $\langle T\rangle \sim M_{\rm Pl}$ returns $\rho \sim M^4_{\rm Pl}$. Neither reproduces $({\rm meV})^4$.

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

$$\rho_\Lambda^{\rm bounce} \sim (\alpha/M)\, M^5_{\rm Pl} \sim 10^{-2}\, M^4_{\rm Pl}, \tag{A2}$$

a bounce-era density that must dilute by $e^{-3N_{\rm tot}} \sim 10^{-122}$ ($N_{\rm tot} \approx 94$) to reach $\rho_\Lambda^{\rm obs}$. Three density symbols appear in what follows and denote three distinct quantities throughout: $\rho_\Lambda^{\rm bounce}$ [Eq. (A2)], $\rho_{\rm bounce}$, and $\rho_\Lambda^{\rm obs}$. 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 $\rho_{\rm bounce} \sim M^4_{\rm Pl}$, 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 $+1 \to +4$ mass-dimension promotion in Case II is a *heuristic dimensional argument*: absorbing the three missing mass powers into on-shell bounce curvature factors ($R \sim M^2_{\rm Pl}$) 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 $\rho_\Lambda^{\rm ECH} \sim M^4_{\rm Pl}$ 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 $M_{\rm Pl}$ (the only scale), so any dimensionally-consistent minimal-ECH parity-odd source sits at

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

Torsion cannot *be* dark energy because dimensional analysis pins its natural density at $M_{\rm Pl}^4$ 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 $+1 \to +4$ dimensional gap, read correctly, is itself the mechanism.

Residual assumption (explicit, not fabricated).—The bound assumes single-scale EFT: $\Lambda \sim M_{\rm Pl}$, no intermediate threshold $\mu \ll M_{\rm Pl}$, and no exact symmetry/topological cancellation. A non-minimal UV completion introducing a new light scale $\mu \ll M_{\rm Pl}$ in the coefficient ($c \sim \mu^a M_{\rm Pl}^{3-a}$), or a symmetry that zeroes the $M_{\rm Pl}^4$ 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 $M_{\rm Pl}^4$ 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 spanning list 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 list; only a non-minimal light scale or protected cancellation—case (ii)—evades it, and that is the tuning being explained.

Inflationary dilution, $\mathcal{D}_{\rm inf} \equiv e^{-3N_{\rm tot}}$, then yields $\rho_\Lambda = \Xi\, M_{\rm Pl}^4$ with $\Xi = (\alpha/M)\, M_{\rm Pl}\, \mathcal{D}_{\rm inf}$ under Eq. (A2). The genuine cosmological-constant hierarchy, evaluated at the exact values fixed above, is $M_{\rm Pl}^4/\rho_\Lambda^{\rm obs} = (1.2209 \times 10^{19}\,\text{GeV})^4/(2.25 \times 10^{-3}\,\text{eV})^4 \approx 8.7 \times 10^{122} \approx 10^{123}$ (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 "$\sim$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 $\rho_{\rm bounce} \sim M_{\rm Pl}^4$, while Eq. (A2) computes the local Case-II pseudo-density $\sim 10^{-2}\, M_{\rm Pl}^4$). The dilution factor required to bridge $\sim 10^{122}$ from $M_{\rm Pl}^4$ down to the observed ρ_Λ is $\mathcal{D}_{\rm inf} \sim e^{-3N_{\rm tot}} \sim 10^{-122}$, giving $N_{\rm tot} \approx 122 \ln 10/3 \approx 94$ e-folds.

Sharper dependency statement.—The precise value of $N_{\rm tot}$ *does* depend on the ansatz choice in Eq. (A2); however, the *overall scale separation* between $\rho_\Lambda^{\rm bounce}$ and $\rho_\Lambda^{\rm obs}$ is $\sim$120 orders of magnitude under *any* non-cancellation assumption (whether the bounce-scale density is $M_{\rm Pl}^4$, $10^{-2}\, M_{\rm Pl}^4$, or $10^{-4}\, M_{\rm Pl}^4$ shifts the required e-fold count by $\mathcal{O}$(a few), not by orders of magnitude), and this scale separation is what drives the 14 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 $N_{\rm tot}$ figure is ansatz-dependent. Readers should therefore treat $N_{\rm tot} \approx 92$–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-list 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 spanning list and show that single-scale NDA closure survives without the Case II dressing.

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

The dimension-+1 operator of Eq. (8) 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–Holst constraint [Eq. (1), of amplitude $\sim \kappa \sim M_{\rm Pl}^{-2}$], and minimally-coupled canonical matter (the construction rule of Sec. V) — and determine the fate of each. As in Sec. V, spanning is asserted from the construction rule rather than proved by exhaustive enumeration, and the list is redundant by construction [Eq. (12)]. No new light scale $\mu \ll M_{\rm Pl}$, no dynamical Immirzi field, and no propagating torsion are admitted, as these lie outside the stated minimal scope. The torsion irrep content is not an assumption but an output of the connection equation: solving it with no restriction on which irreps may appear [Eq. (1)] leaves the tensor (16) irrep identically zero on shell, while the trace-vector (4) irrep is generated by the

Holst term under strictly minimal coupling and is carried throughout. 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 $\varepsilon\, eeR$, Nieh–Yan, and T^2 are component densities of naive dimension $+2$, while Pontryagin $R \wedge R$ and the axial–torsion current structure $\varepsilon\, TeJ^5$ are already dimension $+4$. To write a genuine dimension-4 Lagrangian with *dimensionless* Wilson coefficients we therefore include the single available scale $M_{\rm Pl}^2$ explicitly in each operator definition [Eq. (10)], 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. Two entries survive with nonzero content, and both land in the same Fierz-closed, $M_{\rm Pl}^{-2}$-suppressed four-fermion class. The first is O5's bare $\varepsilon TeJ^5 \;\to\; -3\kappa\,[\gamma^2/(1+\gamma^2)](J^5 \cdot J^5)$ at dim $+4$; only the axial irrep contributes, the trace-vector piece dropping out of the ε contraction. The second is the ε-contracted torsion-square O4. $T_I \wedge T^I$ is supported only by the non-axial torsion irreps: decomposing T_{abc} into its vector, axial, and tensor parts $(4 + 4 + 16 = 24)$, the ε-contracted square vanishes on the pure vector part and on the pure axial part, and is nonzero only on the tensor irrep and on vector$\times$axial cross terms. The on-shell torsion of Eq. (1) has no tensor irrep but a nonzero trace-vector irrep, $\beta/\alpha = 1/(2\gamma)$, so the vector$\times$axial channel is populated and $\mathcal{O}_4^{[4]}$ takes the value of Eq. (13), $-3\kappa\gamma^3/(1+\gamma^2)^2\,(J^5 \cdot J^5)$, whose bare invariant is $-192\pi^2 G^2[\gamma^3/(1+\gamma^2)^2](J^5 \cdot J^5)$. Because that is the same $(J^5 \cdot J^5)$ structure as O5, with $\mathcal{O}_4^{[4]}/\mathcal{O}_5^{[4]} = \gamma/(1+\gamma^2) \simeq 0.22$ at $\gamma = 0.2375$, O4 shares O5's disposal rather than supplying a separate one, and the physics conclusion of this appendix is unchanged. Both statements — the irrep content and the operator values — follow from solving the Holst-modified connection equation itself, varying with respect to all 24 independent contorsion components with no irrep ansatz, cross-checked against an independent differential-form route, and evaluated on six explicitly curved on-shell configurations at $\gamma \in \{19/80, 1, 3\}$ (independent on-shell torsion adjudication module; see the Data and Code Availability statement). The identity $\overline{M}_{\rm Pl}^2\kappa^2 = \kappa$ [exact in the reduced-mass convention $\kappa = \overline{M}_{\rm Pl}^{-2}$ of Sec. II and an order-of-magnitude identity at the full mass] is what converts the bare ε-contracted square into Eq. (13).

Two tensor identities carry the reduction; both are verified symbolically in the released script `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 antisymmetries and the first (algebraic) Bianchi identity $R_{\mu[\nu\rho\sigma]} = 0$,

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

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 branch restriction is not cosmetic: the identity imposed is the *torsion-free* first Bianchi identity, which a nonzero on-shell torsion violates. Off that branch O1 and O6 are not pointwise zero; they are disposed of instead by Eq. (12), which gives $\mathcal{O}_1^{[4]} = \mathcal{O}_6^{[4]} = -\mathcal{O}_2^{[4]} + \frac{1}{2}\mathcal{O}_4^{[4]}$ — an exact total derivative contributing zero to the equations of motion and zero to the vacuum energy, plus the Fierz-closed $M_{\rm Pl}^{-2}$-suppressed contact term of Eq. (13).

Check D — the torsion-square identity, and what it does and does not dispose of.—The minimal (totally antisymmetric) spin current is dual to the axial vector, $S^{abc} = \frac{1}{4}\varepsilon^{abcd}\bar\psi\gamma_d\gamma^5\psi = \frac{1}{4}\varepsilon^{abcd}J^5_d$; using the Lorentzian (mostly-plus, $\varepsilon^{0123} = +1$) contraction $\varepsilon_{abcd}\varepsilon^{abce} = -3!\,\delta^e_d$ (verified symbolically) gives $S_{abc}S^{abc} = \frac{1}{16}\varepsilon_{abcd}\varepsilon^{abce}J^{5d}J^5_e = -\frac{3}{8}\,(J^5 \cdot J^5)$, consistent with the companion paper's torsion-elimination normalization [1], and hence, for the axial part $\alpha\,\varepsilon_{abcd}J^{5d}$ of any torsion, the ε-free square $T_{abc}T^{abc} = -6\alpha^2(J^5 \cdot J^5)$ by the same contraction. That identity concerns the ε-free, parity-*even* square $T_{abc}T^{abc}$. It is a *different* invariant from $\mathcal{O}_4 = \varepsilon^{\mu\nu\rho\sigma}T^I{}_{\mu\nu}T_{I\rho\sigma}$ as defined in Eq. (10), and it does not dispose of O4. The ε-contracted square is instead supported only by the non-axial torsion irreps — it vanishes on the pure vector part and on the pure axial part of T_{abc} and is nonzero only on the tensor irrep and on vector$\times$axial cross terms. The on-shell torsion of Eq. (1) has no tensor irrep, but it does carry a trace-vector irrep, so the vector$\times$axial channel is populated and $\mathcal{O}_4^{[4]}$ is given by Eq. (13) rather than vanishing; in closed form the bare invariant is $-24\alpha\beta(J^5 \cdot J^5)$, zero on a pure axial torsion $(\beta = 0)$ and on a pure trace-vector torsion $(\alpha = 0)$ alike. What Check D does establish, via the same ε-contraction, is the reduction of the linear axial–torsion coupling: $\mathcal{O}_5^{[4]} \to -6\alpha(J^5 \cdot J^5) = -3\kappa[\gamma^2/(1+\gamma^2)](J^5 \cdot J^5)$, a member of the Fierz-closed basis (App. C), bounded uniformly at $M_{\rm Pl}^{-2}$ with a natural *density* scale $\sim M_{\rm Pl}^4$ by single-scale NDA, carried by a dimensionless Wilson coefficient — and O4, at $\mathcal{O}_4^{[4]}/\mathcal{O}_5^{[4]} = \gamma/(1+\gamma^2)$, sits in the same basis at the same power. All of these statements are confirmed by independent exact-rational symbolic computation, off shell and on six explicitly curved on-shell Einstein–Cartan–Holst configurations (independent operator-basis and on-shell torsion adjudication modules; see the Data and Code Availability statement).

Structural entries.—O2 (Nieh–Yan) is an exact 4-form, $d(e_I \wedge T^I) = T_I \wedge T^I - e_I \wedge e_J \wedge R^{IJ}$ [19], and O3 (Pon-

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. (10)] 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. Two entries carry nonzero vacuum-energy content, both inside the Fierz-closed $M_{\rm Pl}^{-2}$-suppressed four-fermion class: O5 at $-3\kappa[\gamma^2/(1+\gamma^2)](J^5\cdot J^5)$ and the ε-contracted torsion-square O4 at $-3\kappa\gamma^3/(1+\gamma^2)^2(J^5\cdot J^5)$ [Eq. (13)], the latter carried by the axial $\times$ trace-vector channel of the on-shell torsion [Eq. (1)]. O1 = O6 is zero on the torsion-free branch and equal to $-$O2 plus $\frac{1}{2}\mathcal{O}_4^{[4]}$ on the on-shell branch [Eq. (12)]. All entries are therefore topological total derivatives ($\to$ zero equations of motion), reduce under the algebraic Cartan–Holst constraint to the Fierz-closed four-fermion basis (App. C; the natural *density* scale set by single-scale NDA is $\sim M_{\rm Pl}^4$, carried by a dimensionless Wilson coefficient), or vanish identically. None yields a (meV)^4 vacuum energy without a new light scale. The list is a spanning list, not a linearly independent basis: it has rank four, with the two exact relations of Eq. (12). The table body abbreviates $(J^5)^2 \equiv J^5\cdot J^5$.

#	dim (bare)	Bare invariant (schematic)	Prefactor	Fate (bare)	Final ($\times$prefactor)
O1	2	$\varepsilon\, e^I e^J R_{IJ}$ (Holst dual)	$M_{\rm Pl}^2$	0 at T=0 (Bianchi, Check A); $-$NY + $\frac{1}{2}\mathcal{O}_4$ on shell	$\frac{1}{2}\mathcal{O}_4^{[4]}$
O2	2	NY $= d(e_I\wedge T^I)$ (Nieh–Yan)	$M_{\rm Pl}^2$	exact 4-form $\to$ 0 EOM	0 (EOM)
O3	4	$R^{IJ}\wedge R_{IJ}$ (Pontryagin)	1	exact $\to$ 0 EOM	0 (EOM)
O4	2	$\varepsilon^{\mu\nu\rho\sigma}T^I{}_{\mu\nu}T_{I\rho\sigma}$ (torsion^2)	$M_{\rm Pl}^2$	$-3\kappa^2\gamma^3(1{+}\gamma^2)^{-2}(J^5)^2$	$-3\kappa\gamma^3(1{+}\gamma^2)^{-2}(J^5)^2$
O5	4	$\varepsilon\, T\, e\, J^5$ (axial–torsion)	1	$\to -3\kappa\gamma^2(1{+}\gamma^2)^{-1}(J^5)^2$, Fierz	$-3\kappa\gamma^2(1{+}\gamma^2)^{-1}(J^5)^2$
O6	2	$\varepsilon^{\mu\nu\rho\sigma}R_{\mu\nu\rho\sigma}$	$M_{\rm Pl}^2$	= O1 exactly (tetrad conversion); same fate	$\frac{1}{2}\mathcal{O}_4^{[4]}$

tryagin) 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, the tensor torsion irrep — 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 (O5, O4, and the non-total-derivative remainder $\frac{1}{2}\mathcal{O}_4^{[4]}$ of O1 = O6 on the on-shell branch) that collapses via the algebraic Cartan–Holst constraint to the Fierz-closed basis at a natural density scale $\sim M_{\rm Pl}^4$; or (iii) a density that vanishes identically — the single-curvature pair (O1, O6) on the torsion-free branch by Eq. (A4). No genuine local dimension-4 parity-odd operator in minimal ECH produces a (meV)^4 vacuum energy without introducing a new light scale $\mu \ll M_{\rm Pl}$ or a protected symmetry — and either *is* the tuning the mechanism was meant to explain. **Single-scale NDA closure therefore survives at genuine dimension 4**, without recourse to the Case II on-shell dressing: the $+1 \to +4$ dressing is dispensable once the true dimension-4 spanning list is exhibited, because every element is topological, Fierz-basis-reducible, or identically vanishing. The dimension-+1 representative of Eq. (8) 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 $\vartheta_{\rm NY}$ a pseudoscalar, $\partial_\mu\vartheta_{\rm NY}$ transforms as a pseudo-co-vector and $J^{5\mu}$ as a pseudo-vector, so their Lorentz-scalar contraction $\partial_\mu\vartheta_{\rm NY}\, J^{5\mu}$ 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 $\langle\partial_\mu\vartheta_{\rm NY}\rangle \neq 0$: 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. (2), 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. (3) proceeds via the standard chiral-anomaly relation $\partial_\mu J^{5\mu} \supset (\alpha_{\rm em}/4\pi)\, F\widetilde{F}$ at the EFT level — the operator of Eq. (2) 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 $\kappa = 8\pi G \sim M_{\rm Pl}^{-2}$, the axial–axial contact term $c_{AA}\,(J^5{\cdot}J^5)$ (the only structure present under *minimal* fermion coupling, per Freidel–Minic–Takeuchi [7]) 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 $M_{\rm Pl}$ power.

On the 16-dimensional Clifford basis, grouped into the five Lorentz classes $(S,V,T,A,P) = (\mathbf{1},\gamma^\mu,\sigma^{\mu\nu},\gamma^\mu\gamma^5,\gamma^5)$, we first state the c-number spinor rearrangement in the normalized convention of Itzykson–Zuber and Nieves–Pal [24, 25]: $e_I(1234) = \sum_J (F_{\rm c})_{IJ}\, e_J(1432)$, with rows I the source classes and columns J the produced classes,

$$F_{\rm c} = \frac{1}{4}\begin{pmatrix} 1 & 1 & \frac{1}{2} & -1 & 1 \\ 4 & -2 & 0 & -2 & -4 \\ 12 & 0 & -2 & 0 & 12 \\ -4 & -2 & 0 & -2 & 4 \\ 1 & -1 & \frac{1}{2} & 1 & 1 \end{pmatrix}, \qquad F_{\rm c}^2 = \mathbb{1}, \tag{C1}$$

in the order (S,V,T,A,P). In this normalized basis $\Gamma^\mu_A = i\gamma^\mu\gamma^5$ 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 $F_{\rm c}$. Turning the c-number identity into an anticommuting field-operator identity requires one Grassmann exchange, $F_{\rm op} = -F_{\rm c}$. 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 $F_{\rm op} = -F_{\rm c}$ yields the operator row $(1,\frac{1}{2},0,\frac{1}{2},-1)$ and the operator-level decomposition

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

while the parity-odd $(J\cdot J^5)$ Holst partner rotates only within the $\{V,A\}$ block (operator-level cross coefficients $F^{\rm op}_{VA} = F^{\rm op}_{AV} = \frac{1}{2}$), 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: $(F_{\rm c})_{AS} = \frac{1}{4}(-4) = -1$, hence $(F_{\rm op})_{AS} = +1$ 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; the elimination chain is carried in App. E) 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 $\gamma\to\infty$ 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 on four *distinct* anticommuting fields, and an exact identical-field Grassmann identity. Uniqueness is a statement about the distinct-field construction only — for a single species the five quartics obey two exact linear relations (span rank three), so identical-field rearrangement rows are *not* unique, and it is the declared direct-channel convention that fixes the mean-field G_s. See the released verification artifact and its report listed in the Data and Code Availability statement. A further released script 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 $\kappa\sim M_{\rm Pl}^{-2}$ 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 $\kappa\sim M_{\rm Pl}^{-2}$ 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 $M_{\rm Pl}^{-2}$-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 $M_{\rm Pl}$-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 $\mu \ll M_{\rm Pl}$, the tensor torsion irrep, 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 $M_{\rm Pl}$-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/$M_{\rm Pl}$; 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 the tensor torsion irrep, which is identically absent from the minimally-coupled on-shell torsion [Eq. (1)] and appears only when the minimal-coupling assumption is relaxed. The trace-vector irrep, by contrast, *is* generated under minimal coupling by the Holst term and is inside the lemma's reach, not outside it: it enters the operator list through O4, whose on-shell value [Eq. (13)] is the same Fierz-closed $(J^5\cdot J^5)$ structure the lemma covers, at the same $M_{\rm Pl}^{-2}$ power. 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. Nothing below is new to this survey. Both halves of the Statement are established here: because Steps 1–4 reduce the action *exactly* to the Einstein–scalar action, the tensor conclusion is an immediate corollary of the scalar one and nothing about it is deferred. What the companion carries in addition, and what this survey does not need, is the explicit order-by-order second-order Holst-term verification and the discussion of what would break the transparency.

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 $e^{[I} \wedge T^{J]} = 0$. In frame components this reads $\delta^{[I}_{[K} T^{J]}_{LM]} = 0$, and for invertible tetrad the linear map it defines on the 24 independent components $T^J{}_{LM}$ ($T^J{}_{LM} = -T^J{}_{ML}$) has trivial kernel, so $T^I = 0$. This is the standard algebraic-torsion property of Einstein–Cartan theory [26]: torsion is not an independent propagating field but is fixed algebraically by — and vanishes with — the spin source. It holds at every classical field configuration in the stated local domain, hence before any perturbative expansion. (3) *The connection reduces to Levi-Civita,* $\Gamma^\lambda{}_{\mu\nu} = \mathring{\Gamma}^\lambda{}_{\mu\nu}$. (4) *The Holst term vanishes by the first Bianchi identity.* Evaluated on the Levi-Civita connection the Holst term reduces to $\frac{1}{2\gamma}\epsilon^{\mu\nu\rho\sigma}R_{\mu\nu\rho\sigma}(\mathring{\Gamma})$ (retaining the explicit γ^{-1} prefactor of Sec. II), 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, stated no more widely than the hypotheses permit: on the zero-spin branch the Holst sector and the Barbero–Immirzi parameter decouple from all scalar and tensor perturbation observables for canonical scalar matter, and in particular minimal ECH produces no helicity-dependent tensor propagation from parity-symmetric initial data. This is B14 (Branch H). Its route reach is R2–R4, where it establishes that the *classical zero-spin* perturbative baseline of each route is exactly inert; it does *not* extend to R1, nor does it imply or confirm B8, because the exclusion list above rules out precisely the nonzero fermion spin density on which both R1 and B8 are built. B8 and B14 are therefore independent statements about the same observable channel on disjoint matter branches, and are counted separately in Table I.

Appendix E: Companion-Imported Tier-II Inputs Carried Self-Contained: Contact Coefficient and R1 Benchmark

Following the same faithful-extraction convention as Appendix D, this appendix carries, from the companion paper's published derivations [1], the two remaining load-bearing companion inputs of this survey: the torsion-elimination chain that fixes the minimal contact operator and its $-(3\kappa/16)[\gamma^2/(1+\gamma^2)]$ coefficient (Sec. II), carried in App. E 1 as Eqs. (E1), (E2) and (E3); and the R1 finite-density benchmark arithmetic quoted in Table II, carried in App. E 2 as Eq. (E5). Nothing below is new to this survey; the derivations follow the companion, which in turn follows Freidel, Minic, and Takeuchi [7].

1. Torsion elimination and the contact coefficient

Define the internal dual on Lorentz bivectors by $(\star X)_{IJ} = \frac{1}{2}\epsilon_{IJ}{}^{KL}X_{KL}$, so $\star^2 = -\mathbb{1}$. Acting on an antisymmetric bivector, the Holst-modified Cartan operator of the first-order ECH action and its inverse are

$$\mathcal{Q}_\gamma = \star + \gamma^{-1}\mathbb{1}, \qquad \mathcal{Q}_\gamma^{-1} = \frac{\gamma^2}{1+\gamma^2}\left(\gamma^{-1}\mathbb{1} - \star\right), \tag{E1}$$

verified by direct multiplication; the inverse exists for every real finite nonzero γ because $1+\gamma^2 > 0$ — the same invertible bivector operator whose triviality of kernel drives Step 2 of Appendix D. Writing the Lorentz-connection source three-form as $\mathcal{J}^{IJ} \propto \delta S_{\rm matter}/\delta\omega_{IJ}$, the connection variation reads $\mathcal{Q}_\gamma(e^{[I} \wedge T^{J]}) = \mathcal{J}^{IJ}$: an algebraic (non-dynamical) equation. For minimal (totally antisymmetric) Dirac coupling, whose spin current is dual to the axial current [$S^{IJK} = \frac{1}{4}\epsilon^{IJKL}J_{5L}$, the Sec. II convention], Freidel, Minic, and Takeuchi solve this equation explicitly [their Eq. (17)], obtaining the contorsion

$$e_I{}^\mu C_{\mu JK} = 4\pi G\,\frac{\gamma^2}{1+\gamma^2}\left(\frac{1}{2}\epsilon_{IJKL}J_5^L - \frac{1}{\gamma}\,\eta_{I[J}J^5_{K]}\right), \tag{E2}$$

Equation (E2) fixes the torsion normalization used throughout this survey. Converted to torsion by $T_{abc} = C_{bac} - C_{cab}$ it is Eq. (1): an axial irrep and a trace-vector irrep in the ratio $\beta/\alpha = 1/(2\gamma)$, and no tensor irrep. The $1/\gamma$ term is that trace-vector piece — its totally antisymmetric part vanishes identically and it carries the entire torsion trace $T^a{}_{ac}$ — not an axial piece written in another basis; being *sourced* by the pseudo-vector J^5 does not place it in the axial irrep, since irrep membership is fixed by index structure. In the $\gamma \to \infty$ Einstein–Cartan limit the trace-vector piece switches off and the torsion is purely axial, $T_{abc} \to \frac{\kappa}{2}\varepsilon_{abcd}J^{5d}$. The direct back-substitution into the action [their Eq. (23)] gives $\mathcal{L}_{\rm int} = -(3/2)\pi G\,[\gamma^2/(1+\gamma^2)]\,J_5^2$ for minimal coupling. The normalization bridge to this survey's $\kappa = 8\pi G$ is

$$4\pi G = \frac{\kappa}{2}, \qquad -\frac{3}{2}\pi G = -\frac{3\kappa}{16}, \tag{E3}$$

so the eliminated-torsion contact operator is

$$\mathcal{L}_{4\psi} = -\frac{3\kappa}{16}\,\frac{\gamma^2}{1+\gamma^2}\,(J_5^I J_{5I}), \tag{E4}$$

exactly the Sec. II operator. The pure Einstein–Cartan coefficient $-(3\kappa/16)$ is recovered for $\gamma \to \infty$; because $\gamma^2/(1+\gamma^2) < 1$ for real finite γ can only reduce the coupling, the companion's gap-equation convention retains the maximal magnitude $G_s = -3\kappa/16$ — the γ-independent channel factor of Appendix C [1]. A vector–axial *four-fermion* term is not generated by minimal coupling; it requires a non-minimal fermion coupling, outside the declared field content.

2. The R1 finite-density benchmark

For a specified homogeneous fermion number density n_ψ the companion defines the coefficient-one dimensional benchmark $\rho_{4f}^{(\rm dim)} \equiv \kappa n_\psi^2 = 8\pi n_\psi^2/M_{\rm Pl}^2$ (mass dimension $+4$, since $[n_\psi] = +3$); the actual axial–axial operator Eq. (E4) adds the factor 3/16. Using $\hbar c = 1.9733 \times 10^{-5}$ eV cm and $M_{\rm Pl} = 1.2209 \times 10^{28}$ eV,

$$\begin{aligned} \kappa n_\psi^2 &\simeq 1.0 \times 10^{-79}\left(\frac{n_\psi}{100\,{\rm cm}^{-3}}\right)^2 {\rm eV}^4, \\ \frac{\kappa n_\psi^2}{\rho_\Lambda} &\simeq 3.6 \times 10^{-69}\left(\frac{n_\psi}{100\,{\rm cm}^{-3}}\right)^2, \end{aligned} \tag{E5}$$

with the companion's normalization $\rho_\Lambda \approx (2.3\,{\rm meV})^4 \approx 2.8 \times 10^{-11}\,{\rm eV}^4$ — about 68.4 orders below ρ_Λ at the deliberately elevated ISM-like normalization $100\,{\rm cm}^{-3}$ (neither a cosmological-density estimate nor a preferred state); including the 3/16 coefficient gives $1.9 \times 10^{-80}\,{\rm eV}^4 = 6.7 \times 10^{-70}\,\rho_\Lambda$. Under the $\rho_\Lambda = (2.25\,{\rm meV})^4$ input of Appendix A the same κn_ψ^2 evaluates to 3.9×10^{-69}, identical at the order-of-magnitude precision quoted (Sec. II). As in the companion, this comparison is a coefficient-scale illustration only: the number density does not determine the state-dependent coincident composite $\langle J^5 \cdot J^5 \rangle$ without specifying the ensemble, spin polarization, and a renormalized contact prescription, and no coherent order parameter, vacuum expectation value, or equation of state is inferred from it.

[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 (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].

[3] D. Benedetti and S. Speziale, Perturbative running of the Immirzi parameter, J. Phys. Conf. Ser. **360**, 012011 (2012), arXiv:1111.0884 [hep-th].

[4] A. Ashtekar and P. Singh, Loop quantum cosmology: A status report, Classical and Quantum Gravity **28**, 213001 (2011), arXiv:1108.0893 [gr-qc].

[5] A. Ghosh and P. Mitra, An improved estimate of black hole entropy in the quantum geometry approach, Phys. Lett. B **616**, 114 (2005), arXiv:gr-qc/0411035 [gr-qc].

[6] S. Holst, Barbero's Hamiltonian derived from a generalized Hilbert-Palatini action, Physical Review D **53**, 5966 (1996), arXiv:gr-qc/9511026 [gr-qc].

[7] 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].

[8] 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].

[9] 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].

[10] D. Benedetti and S. Speziale, Perturbative quantum gravity with the Immirzi parameter, JHEP **2011** (6), 107, arXiv:1104.4028 [hep-th].

[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].

[12] 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].

[13] P. Diego-Palazuelos and E. Komatsu, Cosmic birefringence from the Atacama Cosmology Telescope data release 6, (2025), arXiv preprint; no journal reference at time of writing., arXiv:2509.13654 [astro-ph.CO].

[14] H. Golden, namaster-proof: Exact pseudo-C_ℓ window inference and content-bound validation for reproducible spin-2 analyses, Zenodo 10.5281/zenodo.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.

[15] 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].

[16] 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].

[17] T. Kimura, Divergence of axial-vector current in the gravitational field, Progress of Theoretical Physics **42**, 1191 (1969).

[18] R. Delbourgo and A. Salam, The gravitational correction to PCAC, Physics Letters B **40**, 381 (1972).

[19] H. T. Nieh and M. L. Yan, An identity in riemann–cartan geometry, J. Math. Phys. **23**, 373 (1982).

[20] N. J. Popławski, Cosmology with torsion: An alternative to cosmic inflation, Physics Letters B **694**, 181 (2010), arXiv:1007.0587 [astro-ph.CO].

[21] 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].

[22] I. Brivio and M. Trott, The Standard Model as an Effective Field Theory, Phys. Rept. **793**, 1 (2019), arXiv:1706.08945 [hep-ph].

[23] 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].

[24] C. Itzykson and J.-B. Zuber, *Quantum Field Theory* (McGraw-Hill, 1980).

[25] J. F. Nieves and P. B. Pal, Generalized fierz identities, Am. J. Phys. **72**, 1100 (2004), arXiv:hep-ph/0306087 [hep-ph].

[26] F. W. Hehl, P. von der Heyde, G. D. Kerlick, and J. M. Nester, General relativity with spin and torsion: Foundations and prospects, Reviews of Modern Physics **48**, 393 (1976).