

Channel-Level Amplitude Closure of Four Enumerated Minimal Einstein–Cartan–Holst Dark-Energy Routes Under Stated Assumptions, and Perturbation Transparency for Scalar Matter

Houston Golden^{1, *}

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

(Dated: July 8, 2026)

We assess four enumerated minimal-Einstein-Cartan-Holst (ECH) spin-torsion channels as candidate sources of late-time dark energy and find that each is constrained under stated assumptions: R1 (NJL contact) is *amplitude-suppressed* by a standard torsion-elimination derivation, while R2–R3 (one-loop EA, Immirzi running) are *amplitude-suppressed* under explicitly-labeled scaling ansätze (Sec. IV), and R4 (parity-CMB coupling via spectator ALP or neutrino current) is *not* closed by amplitude mismatch but by an explanatory-deficit / cosmological-constant fine-tuning objection: the same coupling that reproduces β_{obs} requires an ultralight-mass tuning $m_\theta \sim H_0$ to also produce ρ_Λ , relocating the CC problem rather than solving it (Sec. IV F). The phrase “amplitude closure” in the title applies strictly to R1–R3 (which are each ruled out by an explicit amplitude-suppression argument); R4 is *not* amplitude-suppressed but is instead closed by this naturalness/explanatory-deficit route, and is included in the channel-level count because it exhausts the fourth minimal-ECH dark-energy channel. This is a channel-level assessment, *not* an operator-level theorem: the four enumerated routes (NJL, one-loop EA, Immirzi running, parity-CMB) are not proven to be a complete diffeomorphism-invariant operator basis for the minimal-ECH effective action. The two previously-omitted parity-odd operators (Jackiw-Pi gravitational Chern-Simons $R \wedge \tilde{R}$; parity-odd four-fermion partner with $\gamma_{\text{BI}}/(\gamma_{\text{BI}}^2 + 1) \cdot 8\pi G$ coefficient) are now closed explicitly at the operator level (Sec. IV C, Sec. IV B): the former is a total derivative for constant coupling and R4-class otherwise, the latter inherits R1’s M_{Pl}^{-2} suppression as the parity-odd projection of the same torsion-elimination operator. The Fierz-by-Fierz projection lemma establishing that these four-fermion structures close onto a finite, uniformly M_{Pl}^{-2} -suppressed dimension-6 basis is proven in Appendix C; the only residual scope is the non-minimal completion (a new light scale or non-minimal torsion coupling). The dark-energy mapping is a single-scale NDA (naive-dimensional-analysis) dimensional no-go: the parity-odd operator has off-shell mass dimension +1, and single-scale power counting ($\Lambda \sim M_{\text{Pl}}$, no cancellation) forces its natural density to $\rho_\Lambda^{\text{ECH}} \sim M_{\text{Pl}}^4$, never (meV)⁴ (Appendix B). No amplitude is derived, so the argument is not circular; all R4 and dark-energy mapping claims follow from single-scale EFT, the only evasion being a non-minimal light scale or exact cancellation that is itself the tuning being explained. Through 7 foundation studies (Foundations A–G) and 6 observational research branches (Branches H, J, L, M, N, O) we report 13 distinct barriers (Sec. IX; 14 historical catalog entries with B8 subsumed by B14) that collectively constrain the enumerated channels of the minimal-ECH route from the quantum bounce to observable dark energy. These barriers target distinct physical mechanisms (amplitude suppression, thermal washout, operator decoupling, naturalness deficit) and are described as distinct mechanism-class constraints under the classification of Sec. IX. The central result is a perturbation-transparency result: for canonical scalar matter, torsion vanishes at all classical metric/scalar perturbation orders around the torsion-free branch (excluding propagating-torsion, dynamical-Immirzi-field, fermion-loop, and non-minimal-matter sectors), the Holst dual contraction $\epsilon^{\mu\nu\rho\sigma} R_{\mu\nu\rho\sigma}$ vanishes identically on the Levi-Civita connection ($T = 0$) by the first (algebraic) Bianchi identity $R_{\mu[\nu\rho\sigma]} = 0$ (which holds for any torsionless connection), and the Holst sector therefore decouples from all scalar/tensor perturbation equations of motion (Sec. X).^a The proposed link from ECH to a late-time vacuum energy requires a phenomenological dimensional ansatz beyond the minimal framework (Appendix B), and is constrained at the channel-amplitude level by 13 mechanism-class constraints (Sec. IX; 14 historical catalog entries, of which B8 is subsumed by B14 per the perturbation-transparency result), each probing a distinct physical failure mode. A structural tension (Sec. XIV D) exists between the dark-energy mechanism, which requires $N_{\text{tot}} \approx 92$ post-bounce e -folds, and the matter-bounce $f_{\text{NL}} = -35/16$ signature: that many e -folds would *definitively* erase the signature at SPHEREx-accessible comoving wavenumbers. A mode observable today at $k_{\text{SPHEREx}} \sim 10^{-1} h/\text{Mpc}$ maps back to a physical bounce-scale wavenumber $e^{N_{\text{tot}} - N_{\text{exit}}} \sim e^{32}$ larger (for $N_{\text{tot}} \sim 92$, $N_{\text{exit}} \sim 60$), placing it deep inside the inflationary sub-horizon regime, where it carries purely vacuum-inflationary fluctuations rather than matter-bounce contraction modes (the derivation is given in Sec. XIV D). The minimal-ECH four-route channel set is therefore tightly constrained as both a dark-energy generator and a matter-bounce host. The two predictions discussed below as “surviving” are accordingly *not* predictions of ECH itself, but bounce-class and GR+ALP-class observables that the channel-level assessment does not forbid in the parameter regimes complementary to its dark-energy mechanism; we report them here because they remain testable signatures of any bounce model in the broader programme. Specifically: (i) $f_{\text{NL}} = -35/16$ is a property of the *matter-bounce class* [1], derived from the contraction-phase

cubic action with no ECH input (its detection prospects are forecast in a companion [2]; we quote no forecast significance here, as it is not a result of this paper); and (ii) spectator-ALP birefringence $\beta \approx 0.27^\circ$ is a *benchmark consistency point*, not an ECH prediction: it sits inside the WMAP+Planck 1σ band $\beta_{\text{obs}} = 0.342^\circ \pm 0.094^\circ$ ($\sim 3.6\sigma$ from $\beta = 0$, first reported by Minami & Komatsu [3] and refined by Eskilt & Komatsu [4]), and is comparable to the independent ACT DR6 follow-up $\beta = 0.215^\circ \pm 0.074^\circ$ at $\sim 2.9\sigma$ (Diego-Palazuelos & Komatsu [5]; these two birefringence significances arise from different null procedures and are not directly comparable in a single tension table); the same benchmark arises in any GR+ALP setup with the same parameters and is *not* derived from the ECH action. The role of this paper is thus the *channel-level closure* of the four enumerated minimal-ECH dark-energy routes (Sec. IV) at amplitude-budget granularity, basis-complete within minimal ECH at the M_{Pl} -power-counting level as stated above. The complementary observational programme— $\Lambda\text{CDM} + \Delta N_{\text{eff}}$ MCMC verification, NaMaster pipeline validation, and ALP parameter fitting—is documented in the appendices of this paper (Appendices E–H); all reproducibility artifacts are archived with this submission.

PACS numbers: 98.80.-k, 04.50.Kd, 04.60.Pp, 95.36.+x

CONTENTS

I. Introduction	3	B. Parity-odd four-fermion Holst partner of R1: closed by the same Planck suppression and vanishing mean field	14
A. Theoretical Foundations and Novel Synthesis	4	C. Jackiw–Pi gravitational Chern–Simons $R \wedge \tilde{R}$: closed as a total derivative for constant coupling; R4-class otherwise	15
B. Paper Organization	5	D. Route 2 (one-loop graviton corrections to the Holst sector): closed by parity-odd coefficient and Planck suppression	15
II. Theoretical Framework	7	E. Route 3 (quantum running of the Immirzi parameter): closed by mass-dimension lock	17
A. Loop Quantum Cosmology and the Holst Action	7	F. Route 4 (parity-odd CMB coupling via spectator ALP or neutrino current): naturalness objection rather than amplitude no-go	18
1. Einstein-Cartan-Holst Action	7	G. Closure summary	19
2. Derivation of the Parity-Odd Term	7	V. Data Methods: Galaxy Spin Analysis	20
3. Parameter Naturalness	9	VI. Systematic Analysis	20
B. Black Hole Interior and Quantum Bounce	9	VII. Falsifiability Criteria	20
C. Cosmic Rotation and Dark Energy	9	VIII. Related Work	21
1. Inflationary Suppression	10	IX. Structural Constraints on Dark-Energy Routes in Minimal ECH	22
2. Galaxy Spin Alignment Mechanism	11	A. Barrier 1: Mass-Coupling Lock (Foundation A)	22
III. Observational Signatures	11	B. Barrier 2: Topological-Shift Duality (Foundation B)	22
A. CMB E - B Cross-Correlations	11	C. Barrier 3: Scalar-Tensor Universality (Foundation C)	22
B. Galaxy Spin Asymmetry: A Confirmed Null	12	D. Barrier 4: Planck Suppression (Foundation D)	22
IV. Four-Route No-Go: Why Each Standard ECH Channel Closes	12	E. Barrier 5: Scale Separation (Foundation E)	23
A. Route 1 (NJL four-fermion contact): closed by Planck suppression	14	F. Barrier 6: Attractor-Sensitivity Dilemma (Foundation F)	23
		G. Barrier 7: Parameter Immunity (Foundation G)	23

* houston@hubify.com

^a This Bianchi-identity vanishing is distinct from — and should not be confused with — the Pontryagin density $\propto R \tilde{R} = \epsilon^{\mu\nu\rho\sigma} R_{\mu\nu}{}^{\alpha\beta} R_{\rho\sigma\alpha\beta}$, which involves *two* curvature tensors and is a separate topological invariant. The Holst dual contraction has only *one* curvature. In differential-form language, with the standard Nieh–Yan density $\text{NY} \equiv d(e_I \wedge T^I) = T_I \wedge T^I - e_I \wedge e_J \wedge R^{IJ}$, one has $e^I \wedge e^J \wedge R_{IJ} = -\text{NY} + T^I \wedge T_I$ (Nieh–Yan density plus torsion-squared); both pieces vanish at $T = 0$, recovering the Bianchi-vanishing result. The headline conclusion (Holst sector decouples from scalar/tensor EOM) follows from the Bianchi-trivial argument, not from a Pontryagin total-derivative argument.

H. Barrier 8: Parity-Even Interaction (Branch H)	23	E. Cosmological Constraint Methodology: Λ CDM+ ΔN_{eff} MCMC Proxy and the ECH-Sector ΔN_{eff} Bound	34
I. Barrier 9: Liouville Conservation (Branch J)	23	1. Cosmological Tensions: H_0 and σ_8	34
J. Barrier 10: UV \rightarrow IR Specificity Dilemma (Branch L)	23	2. Stock-CAMB Λ CDM+ ΔN_{eff} MCMC: Generic Radiation-Proxy Test (Not a Spin-Torsion Theory Module)	34
K. Barrier 11: Decoupling Universality (Branches L/M)	24	a. Bespoke ECH-sector ΔN_{eff} from the torsion-induced four-fermion interaction	34
L. Barrier 12: Vacuum Amplification Ceiling (Branch M)	24	3. Cosmological Fits and Model Comparison	38
M. Barrier 13: Gravitational Democracy (Branches N/O)	24	a. Datasets and Configuration	38
N. Barrier 14: Perturbation Transparency	24	b. Results: Λ CDM+ ΔN_{eff} proxy	38
X. The Perturbation-Transparency Result	24	F. NaMaster Pseudo- C_ℓ $E \rightarrow B$ Pipeline Validation	38
A. Statement	24	1. Data Methods: CMB E - B Analysis	38
B. Proof (Scalar Sector)	24	G. Spectator-ALP Cosmic-Birefringence Consistency Check	42
C. Extension to Tensor Sector	25	1. Cosmic Birefringence: Spectator ALP Consistency Check	42
D. Explicit Verification: The Holst Term in Perturbation Theory	25	ALP dark-energy fraction Ω_a : definition and computation	45
E. What Would Break the Transparency	25	H. Data Availability, Reproducibility Materials, and Sampled-Parameter Priors	46
F. Implications	26	1. Reproducibility Materials	46
G. Discrimination Among Bouncing Cosmologies	26	2. Claims Classification	47
XI. The Hybrid Dark-Energy Loophole	26	3. ALP-MCMC Sampled Parameters, Priors, and Likelihood Stack	47
XII. Discussion	26	References	49
A. The Inflationary Suppression Factor	26		
B. Theoretical Implications	27		
XIII. Surviving ECH-Independent Class Tests	28		
XIV. Limitations and Future Directions	28		
A. Current Limitations	28		
1. Theoretical	28		
2. Observational	28		
B. Robustness to Galaxy Spin Null Results	29		
C. Discriminating Observational Channels	29		
D. Structural Tension: Dark Energy vs. Bounce f_{NL}	29		
E. Channel-Level Closure	29		
XV. Conclusions	29		
Data and Code Availability	30		
Acknowledgments	30		
A. Complete Parameter Summary	31		
B. Dimensional Status of the Parity-Odd Operator	31		
C. Fierz-by-Fierz Projection Lemma for the Minimal-ECH Four-Fermion Basis	32		
D. Line-of-Sight Birefringence from the Maxwell-Chern-Simons Operator	33		

I. INTRODUCTION

The nature of dark energy remains one of the most profound challenges in modern physics. While the Λ CDM model successfully accounts for observed cosmic acceleration [6], it faces severe theoretical difficulties—most notably the cosmological constant problem [7]. DESI 2024–2025 BAO results suggest dynamical dark energy at 3.1 – 4.2σ (dataset-dependent) [8, 9], adding urgency to the search for extensions of the standard model.

This work presents and closes a phenomenological framework that connects ECH spin-torsion cosmology to dark energy. The structural conclusion is a *channel-level amplitude closure* of the four enumerated minimal-ECH dark-energy routes (Sec. IV): under the stated assumptions, the 14 constraints (Sec. IX; 13 distinct mechanism-class constraints, with B8 subsumed by B14) close those routes at amplitude-budget granularity. The surviving phenomenological predictors—matter-bounce f_{NL} and spectator-ALP birefringence—are ECH-independent class tests (shared with other UV completions within the broader bounce/ALP landscape; see §XIII for the precise scoping).

a. Scope and limitations. The four routes assessed below are minimal channels; within the minimal ECH field content they are shown to be basis-complete at the level of M_{Pl} -power-counting classes (§II C, completeness lemma). In particular, the Jackiw–Pi gravitational Chern–Simons term $R\tilde{\wedge}\tilde{R}$ and the parity-odd four-fermion partner of Route 1 (carrying coefficient $\gamma_{\text{BI}}/(\gamma_{\text{BI}}^2+1)\cdot 8\pi G$)—previously flagged as omitted operators—are now closed explicitly in the body (§II C), covered by the completeness argument, now made explicit by the Fierz-by-Fierz projection lemma of Appendix C. The dark-energy mapping presented in Sec. II C and Appendix B is a single-scale NDA dimensional no-go: the leading parity-odd operator written in Eq. (6) has off-shell mass dimension +1, and single-scale power counting ($\Lambda \sim M_{\text{Pl}}$, no intermediate threshold, no cancellation) forces its natural density to $\rho_{\Lambda}^{\text{ECH}} \sim M_{\text{Pl}}^4$, never $(\text{meV})^4$ (see Appendix B for the two admissible completions). No positive amplitude is derived, so the argument is not circular; all R4 and dark-energy mapping claims in this paper follow from single-scale EFT power counting, with the only evasion being a non-minimal light scale or exact cancellation that is itself the tuning being explained. The perturbation-transparency result of Sec. X is restricted to canonical scalar matter: fermion spin density, propagating torsion (Poincaré gauge theory), dynamical Immirzi fields, non-minimal matter couplings, and boundary/topological sectors are explicitly outside its scope. Within those caveats, the four channels are assessed at the amplitude-budget granularity at which observations discriminate, and the resulting closure is a channel-level statement under specified assumptions rather than a full operator-level theorem.

b. Self-containment and companion dependency. The logical content of this paper—the perturbation-transparency theorem (Sec. X), the four-route channel-level closures (Sec. IV), and the 14-entry constraint catalog (Sec. IX)—is established analytically within this manuscript and does *not* depend on any companion result. The companion works (the folded MCMC/reproducibility appendices, II, III, IV; all posted concurrently in the coordinated submission, which fixes only the citable arXiv identifiers) are cited only to anchor *illustrative* numerical values—MCMC posteriors, Fisher-forecast significances, and pipeline validations summarized in Table II—that contextualize the broader observational programme; *none of these companion-imported numbers is load-bearing for any closure, no-go, or theorem stated here.* A referee unable to access the companion papers can therefore still audit every primary claim directly from this paper’s own equations and assumptions. Moreover, the underlying computational artifacts behind those illustrative numbers—the frozen MCMC chains, the NaMaster pseudo- C_{ℓ} validation suite, the spectator-ALP fit, and the galaxy/anomaly catalogs—are archived with this submission in the public reproducibility tree [10], so they can be inspected and re-derived *now*, independently of the companions’ posting

What this paper does and does not establish.

Establishes (within stated scope): (i) a rigorous perturbation-transparency result—for canonical scalar matter the Holst sector decouples from all scalar/tensor perturbation equations of motion (Tier I; Sec. X); (ii) a *channel-level amplitude assessment* of four *enumerated* minimal-ECH dark-energy routes (R1–R4), finding each constrained: R1 parity-even and M_{Pl}^{-2} -suppressed (structural), R2–R3 amplitude-suppressed *under explicitly-labeled scaling ansätze* (Tier III), R4 closed by a naturalness / explanatory-deficit objection *not* an amplitude exclusion (Tier II).

Does not establish: (a) a fully general operator-level no-go—the four routes are not proven a complete diffeomorphism-invariant operator basis; within *minimal* ECH they are basis-complete at the M_{Pl} -power-counting level (the Jackiw–Pi $R\tilde{\wedge}\tilde{R}$ term and the parity-odd four-fermion Holst partner, previously flagged as omitted, are closed explicitly in §II C), with the fully explicit Fierz-by-Fierz projection lemma proven in Appendix C; (b) a first-principles positive-amplitude ECH-to- ρ_{Λ} derivation—none is claimed; the dark-energy mapping is instead a single-scale NDA dimensional no-go (off-shell mass dimension +1 forces $\rho_{\Lambda}^{\text{ECH}} \sim M_{\text{Pl}}^4$, never $(\text{meV})^4$; App. B), and every R4 / dark-energy claim is conditional on single-scale EFT ($\Lambda \sim M_{\text{Pl}}$, no cancellation); (c) any claim resting on the companion papers—their illustrative numbers are non-load-bearing for the results above.

The phrase “four-route closure” throughout means exactly (ii): a channel-level, assumption-conditional amplitude statement, *not* an operator-level theorem.

schedule.

Our approach builds on three theoretical pillars:

1. *Loop Quantum Cosmology* (LQC), providing a non-singular quantum bounce replacing the classical Big Bang singularity [11]. The bounce occurs at $\rho_{\text{crit}} \simeq 0.27\text{--}0.41 \rho_{\text{Pl}}$ (Barbero-Immirzi entropy-counting scheme dependent; Sec. II B).
2. *Einstein-Cartan theory* incorporating fermionic spin-torsion coupling, generating four-fermion contact interactions and preventing gravitational singularities through torsion-induced repulsion at extreme densities [12, 13].
3. *Black hole universe origin*, where a rotating parent black hole spawns a non-singular baby universe through torsion-regulated gravitational collapse [14]. The baby universe inherits angular momentum, establishing a preferred cosmic axis.

A. Theoretical Foundations and Novel Synthesis

Our framework collects well-established theoretical components and tests them as a channel-level amplitude closure of the four enumerated minimal-ECH dark-energy

TABLE I. Executive summary: systematic investigation of bounce-cosmology dark-energy routes within ECH. Structural constraints narrow phenomenological pathways. The $f_{\text{NL}} = -35/16$ matter-bounce signature [1] is a testable prediction of the broader matter-bounce *class*, *not* of the ECH dark-energy mechanism: as Sec. XIV D proves, the $N_{\text{tot}} \approx 92$ e -folds required to source ρ_{Λ} would erase this signature at SPHEREx scales, so the dark-energy route and the observable matter-bounce program are *mutually exclusive*. Quantitative cosmological values cited below are companion-MCMC inputs (Table II), not results of this paper.

Question	Result	Status
Can bounce derive dark energy?	14 constraints map minimal-ECH route space	Phen. assumption ^a required.
Is there a nonsingular bounce?	LQC: $\rho_c \simeq 0.27\text{--}0.41 \rho_{\text{Pl}}$	Yes (LQC holonomy).
ECH visible in scalar/tensor pert.?	Perturbation-transparency result	Decouples; tests \rightarrow ALP/GW.
Testable prediction?	$f_{\text{NL}} = -35/16$ (matter-bounce class ^b)	Class-level ^c ; <i>erased if ECH-DE^d</i> .
Mechanism-independence?	Class-level: scalar-only $w = 0$ matter-bounce ^c	Not a distinctive ECH prediction.
H_0/σ_8 tension resolution?	$\Delta N_{\text{eff}} \approx 0$ (companion MCMC ^e)	Recovers Λ CDM.

^aReparameterized as sensitivity to N_{tot} ; not solved. ^b $f_{\text{NL}} = -35/16$ is the analytic matter-bounce contraction-phase value [1], derived with no ECH input; it is in principle accessible to a multi-tracer SPHEREx bispectrum analysis. The detailed Fisher significance is forecast in a companion (companion [2]) and is *not* quoted here, as it is not a result of the present paper.

^cClass-level: scalar-only $w = 0$ matter-bounce under Assumption (f) of [2]; not fully mechanism-independent across the bouncing-cosmology landscape; not a distinctive ECH prediction. ^dThe dark-energy route requires $N_{\text{tot}} \approx 92$ post-bounce e -folds, which pushes the matter-bounce signature deep into the inflationary subhorizon regime and *definitively erases* it at SPHEREx-accessible scales (Sec. XIV D). The two channels cannot both be realized. ^eCompanion-MCMC input ($H_0 = 67.68 \pm 1.06$, $\Delta N_{\text{eff}} = -0.020 \pm 0.169$; Appendix E, coordinated submission), consolidated in Table II; not load-bearing for any closure or theorem here.

routes (we do not claim a full operator-basis closure; see Sec. IV “Scope” for the explicit operators omitted). The original contributions are:

1. *14-constraint catalog and perturbation-transparency result*: Through 7 foundation studies (Foundations A–G) and 6 observational research branches (Branches H, J, L, M, N, O), we establish 14 mechanism-class structural constraints (one of which, B8, is the observational consequence of the perturbation-transparency result B14 and is retained in the catalog for historical mechanism-class completeness) mapping the minimal ECH parameter space. The central result is that minimal ECH gravity is perturbation-transparent for canonical scalar matter: torsion vanishes at all orders, the Holst sector decouples cleanly from scalar/tensor observables, and parity-sensitive channels (model-dependent tests of γ_{BI} only under a derived γ_{BI} -dependent photon or tensor-parity coupling) shift to nonperturbative parity-violating channels (ALP birefringence, primordial GWs).
2. *Structural tension between dark-energy suppression and bounce f_{NL}* : We identify an incompatibility between the inflationary-suppression dark-energy mechanism ($N_{\text{tot}} \approx 92$ e -folds required) and the matter-bounce $f_{\text{NL}} = -35/16$ prediction (*definitively* erased once the bounce-vs-CMB-horizon-exit differential $N_{\text{tot}} - N_{\text{exit}}$ exceeds the matter-bounce contraction-mode coherence window $N_{\text{coh}} \sim \mathcal{O}(\text{few})$, since the SPHEREx accessible band $k \sim 10^{-4}\text{--}10^{-1} h/\text{Mpc}$ maps to bounce-era *physical* scales $k_{\text{bounce}}^{\text{phys}} = k_{\text{SPHEREx}}^{\text{phys}} e^{N_{\text{tot}} - N_{\text{exit}}}$ at $N_{\text{tot}} \sim 92$ and $N_{\text{exit}} \sim 60$ (relative e -fold differential ~ 32 , deeply inside the erasure regime; comoving k are constant

by definition, only physical scales scale with $a^{-1} \propto e^{-N}$) that lie deep inside the inflationary subhorizon regime where the observable bispectrum is dominated by vacuum-inflationary modes rather than the matter-bounce contraction modes; see Sec. XIV D), establishing that these are independent observational programs.

3. *Survival of ECH-independent class tests*: Despite ECH structural closure, two class-level predictions of the broader bounce/ALP landscape survive and are testable by SPHEREx (2028) and LiteBIRD (early 2030s) independently of the specific ECH UV completion.

B. Paper Organization

Section II develops the ECH theoretical framework. Section III presents observational signatures (galaxy spin null, CMB EB). Section IV closes each of the four standard ECH routes (NJL, one-loop EA, Immirzi running, parity-CMB) at the amplitude level under explicitly-labeled scaling assumptions for R2–R3 and a naturalness limit for R4 (see Sec. IV for the per-route scoping). Section V covers the galaxy-spin data methods. Sections VI–VIII provide systematic analysis, falsifiability criteria, and related work. The core results occupy Secs. IX–XI: the 13 mechanism-class structural barriers (14 historical catalog entries), perturbation-transparency result, and hybrid loophole rejection. Section XII discusses the inflationary suppression factor and theoretical implications. Section XIII summarizes surviving ECH-independent class tests. Section XIV addresses limitations. Section XV concludes.

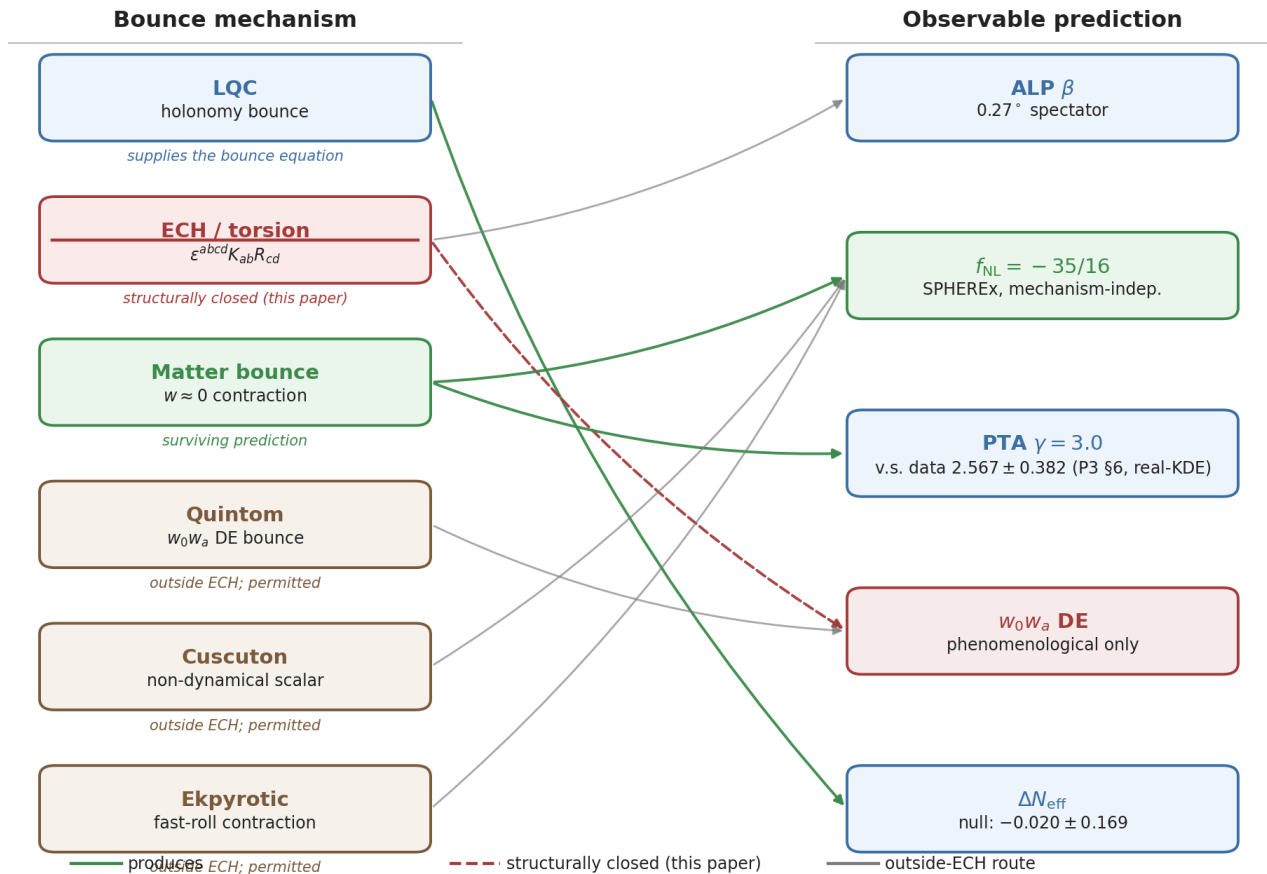


FIG. 1. Bounce-mechanism \rightarrow observable-prediction map. Left column: candidate non-singular bounce mechanisms (LQC, ECH/torsion, matter bounce, quintom-B, Cuscuton, ekpyrotic). Right column: distinctive observable channels. ECH appears bordered with a dashed box marked *channel-level closure under stated assumptions (this paper)*—the 14-constraint catalog narrows the four enumerated minimal-ECH dark-energy channels to zero phenomenologically free pathways within those channels. The PTA annotation reflects the current real-KDE reanalysis $\gamma_{\text{PTA}} = 2.567 \pm 0.382$ (Sec. X G); γ_{PTA} denotes the GWB power-law spectral index, distinct from the Barbero-Immirzi parameter γ throughout this paper.

Companion paper.— $\Lambda\text{CDM} + \Delta N_{\text{eff}}$ MCMC verification (Cobaya v3.6.1, **309,189** frozen accepted samples across two converged dataset combinations: 176,240 full-tension + 132,949 Planck+BAO+SN; see Appendix E Table I for the per-dataset breakdown), NaMaster pseudo- C_ℓ pipeline validation, spectator-ALP MCMC parameter fitting, and a cross-paper reproducibility manifest are reported in Appendix E. Cosmological parameter values referenced in this paper ($H_0 = 67.68 \pm 1.06$, $\Delta N_{\text{eff}} \approx 0$, etc.) are drawn from the companion internal MCMC analysis (Appendix E, *companion; coordinated submission*); they are documented internally rather than as externally citable arXiv-posted numbers, and should be read as internal-analysis inputs to the present structural argument rather than as independently peer-reviewable values until the folded MCMC/reproducibility appendices are publicly posted. We emphasize: *none of these companion-imported numerical values is used in the channel-level closure proof of Sec. IV or the 13*

mechanism-class structural barriers of Sec. IX; the structural closure rests on the dimensional / operator-counting / perturbation-transparency arguments alone, which are evaluable without the companion’s MCMC posteriors. The companion values are quoted only to anchor the discussion of ΛCDM -consistency in Sec. III and the parameter-summary table in Appendix A. All subsequent citations to Appendix E in this paper refer to this companion work; details are not duplicated here. Table II consolidates every companion-imported quantity referenced anywhere in this paper—its value, the companion it is drawn from, and a one-line statement of how it was obtained—so that a referee can audit the present paper’s logic without the companion papers in hand; these inputs are summarized here for self-contained readability, while their full derivations reside in the companion papers posted concurrently in the coordinated submission. The distinction is sharp: the no-go logic itself (the perturbation-transparency theorem, the four-route

channel-level closures, and the mechanism-class constraint catalog) is self-contained and uses *none* of these numbers; the imported values enter only the *illustrative* observational comparisons (the Λ CDM-consistency check of Sec. III and the surviving-signature context of Sec. XIII).

Appendices provide the parameter summary (Appendix A) and dimensional analysis (Appendix B). Supplementary materials are at <https://github.com/Hubify-Projects/bigbounce>.

II. THEORETICAL FRAMEWORK

A. Loop Quantum Cosmology and the Holst Action

1. Einstein-Cartan-Holst Action

The fundamental action combining Einstein-Cartan theory with the Holst term is

$$S_{\text{ECH}} = \frac{1}{16\pi G} \int d^4x e \left[e_a^\mu e_b^\nu R_{\mu\nu}^{ab} + \frac{1}{\gamma} \varepsilon^{abcd} e_a^\mu e_b^\nu R_{cd\mu\nu} + \frac{1}{4} T^{abc} T_{abc} \right] + S_{\text{matter}}, \quad (1)$$

where $e = \det(e_\mu^a)$ is the tetrad determinant, $R_{\mu\nu}^{ab}$ is the curvature of the Lorentz connection, γ is the Barbero-Immirzi parameter, and T^{abc} is the torsion tensor.¹ The Holst term contributes non-trivially when fermions are present, as established by Freidel, Minic & Takeuchi [16], who showed that the Barbero-Immirzi parameter becomes physically observable through its coupling to fermionic matter; the explicit γ -dependence it induces appears in the four-fermion contact term, Eq. (4). To state the variational principle unambiguously: Eq. (1) is a *first-order* (Palatini–Einstein–Cartan) action whose two independently varied fields are the vielbein e_μ^a and the Lorentz spin connection ω_μ^{ab} (equivalently its torsion part), together with the Dirac field ψ in S_{matter} ; the metric-compatible Levi-Civita piece is not a separate field. The connection is algebraic (non-propagating): its

equation of motion, Eq. (3), is the purely algebraic Cartan constraint $T^{abc} = \kappa S^{abc}$, solved pointwise and back-substituted. The $T^{abc} T_{abc}$ term in Eq. (1) is a shorthand for the four-fermion contact interaction obtained after integrating out the non-propagating torsion; it is *not* one of the independently varied fields and is *not* varied on its own — the full variation is performed with respect to $\{e_\mu^a, \omega_\mu^{ab}, \psi\}$ on the Einstein–Cartan–Holst+Dirac action, with Eq. (3) the resulting Cartan equation and $\frac{1}{4} T^{abc} T_{abc}$ appearing only *after* on-shell torsion elimination, so the principle is well-defined and no double counting arises.

The Barbero-Immirzi parameter is fixed by LQG black hole entropy in the SU(2) full-counting scheme used here:

$$\gamma_{\text{SU}(2)} \approx 0.274, \quad (2)$$

where the apparent uncertainty range is *scheme dependence rather than a statistical or theoretical error*: the U(1) horizon-state counting [17] gives $\gamma_{\text{U}(1)} \approx 0.127$ (using $\gamma = \ln 2 / (\pi\sqrt{3})$), the refined SU(2) full counting [18, 19] gives $\gamma_{\text{SU}(2)} \approx 0.274$ (adopted in this paper), and the further Domagała–Lewandowski–Meissner refinement gives $\gamma_{\text{DLM}} \approx 0.2375$. Domagała–Lewandowski and Meissner do *not* quote a ± 0.020 statistical uncertainty; the ~ 0.037 figure that appears in the parameter-budget table (Appendix A) is the spread *between* counting prescriptions (specifically the SU(2)–DLM pair: $\gamma_{\text{SU}(2)} \approx 0.274$ minus $\gamma_{\text{DLM}} \approx 0.2375$ gives $0.0365 \approx 0.037$), retained as an effective range only and *not* propagated as a statistical error in any quantitative claim.

2. Derivation of the Parity-Odd Term

Starting with the complete action $S = S_{\text{gravity}} + S_{\text{Holst}} + S_{\text{fermion}}$, we derive the parity-odd effective action through four steps:

Step 1: Torsion Activation.—Torsion is determined algebraically by the fermionic spin density:

$$T^{abc} = 8\pi G S^{abc}, \quad (3)$$

where $S^{abc} = \frac{1}{4} \bar{\psi} \gamma^{[a} \gamma^{bc]} \psi$.²

¹ Two convention notes for Eq. (1). (i) The displayed $\frac{1}{4} T^{abc} T_{abc}$ is an on-shell Hehl–Datta *shorthand* for the four-fermion contact term obtained after eliminating the non-propagating torsion via the Cartan equation Eq. (3); it is not an independent kinetic term and is *not* varied independently (see the elaboration immediately below and the single-convention footnote at Eq. (3)). (ii) Placing the overall $1/(16\pi G)$ outside the bracket as displayed gives a Holst prefactor $1/\gamma$ on the $\varepsilon^{abcd} e_a^\mu e_b^\nu R_{cd\mu\nu}$ contraction; the equivalent form-language expression $\frac{1}{2\kappa} \int e \wedge e \wedge (R + \frac{1}{\gamma} \star R)$ used by [15, 16] carries $\frac{1}{2\gamma}$ on the Holst $\star R$ term because the $e \wedge e$ wedge product supplies an additional antisymmetrization factor of $1/2$; the two component-form conventions are numerically identical.

² Single-convention derivation (fixing the paper’s normalization once): we use the form-language torsion definition $T^a \equiv De^a$ (components $T^\lambda_{\mu\nu} = 2\Gamma^\lambda_{[\mu\nu]}$) and the Hermitian (symmetrized) minimally coupled Dirac action. Varying that action with respect to the connection gives the spin current $S^{abc} \equiv \delta\mathcal{L}_\psi / \delta\omega_{abc} = \frac{1}{4} \bar{\psi} \gamma^{[a} \gamma^{bc]} \psi = \frac{1}{4} \varepsilon^{abcd} J_d^5$ — *totally antisymmetric*, so all torsion trace parts vanish and the Cartan equation is exactly $T^{abc} = \kappa S^{abc}$, Eq. (3). Consistency proof by back-substitution: $S_{abc} S^{abc} = \frac{1}{16} \varepsilon_{abcd} \varepsilon^{abce} J_e^5 J_e^5 = -\frac{3}{8} J_\mu^5 J^{5\mu}$ (Lorentzian $\varepsilon_{abcd} \varepsilon^{abce} = -3! \delta_d^e$, with mostly-plus signature $g = \text{diag}(-, +, +, +)$ and $\varepsilon^{0123} = +1$); explicitly, the displayed $\frac{1}{4} T^{abc} T_{abc}$ in Eq. (1) is not an independent kinetic term but the on-shell Hehl–Datta *shorthand* for the four-fermion con-

TABLE II. Companion-imported inputs referenced in this paper, summarized here for self-contained readability. For each quantity we give the value, the companion it is drawn from, and a one-line statement of how it was obtained. These values provide *illustrative* observational context only: *none* is load-bearing for the perturbation-transparency theorem (Sec. X), the four-route channel-level closures (Sec. IV), or the mechanism-class constraint catalog (Sec. IX), all of which are established analytically within this manuscript. Full derivations are in the named companions, posted concurrently in the coordinated submission; numerical values are reproduced from Table VI. The underlying computational artifacts are archived with this submission and are refereeable now, independently of the companions’ posting schedule [10]: the $H_0/\Delta N_{\text{eff}}/\sigma_8/\Omega_m$ posteriors and the spectator-ALP fit from the frozen Cobaya chains (`reproducibility/cosmology/chains`), the NaMaster *EB* validation (`reproducibility/p1_namaster_500mc`), and the galaxy-chirality and multi-survey anomaly catalogs (`pipelines/p2_chirality`, `pipelines/p3_anomaly_engine`). External (non-companion) measurements are cited inline at point of use, not here.

Quantity	Value	Companion	How obtained
H_0	67.68 ± 1.06 km/s/Mpc	Appendix E	Joint Λ CDM+ ΔN_{eff} MCMC posterior; Cobaya v3.6.1 + CAMB on Planck+BAO+SN (309,189 frozen accepted samples).
ΔN_{eff}	-0.020 ± 0.169 (full-tension)	Appendix E	Same joint MCMC chain; consistent with 0.
σ_8	0.803 ± 0.008	Appendix E	Derived parameter of the same MCMC chain.
Ω_m	0.308 ± 0.005	Appendix E	Derived parameter of the same MCMC chain.
SPHEREx f_{NL} forecast	1.3–2.75 σ realistic (2.6–2.75 σ optimistic)	Paper II [2]	Multi-tracer SPHEREx Fisher forecast recasting Heinrich+2024 $\sigma(f_{\text{NL}}) \approx 0.7$ (ideal) / ≈ 1.0 (degraded-with-systematics) for the $f_{\text{NL}} = -35/16$ template.
Spectator-ALP β benchmark	$\approx 0.27^\circ$	Appendix E	Spectator-ALP MCMC fit (9,720 accepted samples, $\hat{R} - 1 < 0.01$) at $f_a \sim M_{\text{Pl}}$, $m \sim H_0$; sits inside the external WMAP+Planck band $0.342^\circ \pm 0.094^\circ$ [4].
NaMaster <i>EB</i> pipeline	validated	Appendix E	Pseudo- C_ℓ (<i>EB</i>) estimator validation against simulations; methodology and reproducibility manifest documented in the companion.

Step 2: Four-Fermion Contact Interaction.— Substituting Eq. (3) and integrating out torsion:

$$\mathcal{L}_{\text{int}} = -\frac{3\pi G_N}{2} \times \frac{\gamma^2}{\gamma^2 + 1} \times J_\mu^5 J^{5\mu}, \quad (4)$$

where $J^{5\mu} = \bar{\psi} \gamma^\mu \gamma^5 \psi$ is the axial current, and G_N denotes Newton’s gravitational constant (identified with

tact interaction obtained after eliminating the algebraic torsion via Eq. (3). Performing the connection variation on the full Einstein–Cartan–Holst+Dirac action and integrating out the non-propagating torsion (Hehl 1976 [12]; Freidel–Minic–Takeuchi 2005 [16]) yields the single net contact term $\mathcal{L}_{\text{int}} = -\frac{3\kappa}{16} J_\mu^5 J^{5\mu}$, which is the $\gamma \rightarrow \infty$ (pure Einstein–Cartan) limit of the displayed Eq. (4) and the Hehl–Datta coefficient of Eq. (13). The intermediate algebraic cancellation between the gravitational $T \cdot T$ piece (mass dimension $[\kappa^2] S \cdot S$) and the fermionic spin-connection coupling (linear in κ) drops one power of κ ; we do not reproduce that algebra here and refer the reader to Hehl 1976 Eq. (3.20)–(3.21) and Freidel–Minic–Takeuchi 2005 Eqs. (7)–(13) for the single-step derivation. The $T^{abc} = (\kappa/2) \bar{\psi} \gamma^{[a} \gamma^b \gamma^c] \psi$ form quoted in the original Hehl–Datta-era literature (Sec. IV A) uses the half-weight torsion definition $T^\lambda{}_{\mu\nu} = \Gamma^\lambda{}_{[\mu\nu]}$ and maps to Eq. (3) exactly (the full-weight definition $T^\lambda{}_{\mu\nu} = 2\Gamma^\lambda{}_{[\mu\nu]}$ used here absorbs the factor of 2, so the half-weight ($\kappa/2$) prefactor and the full-weight κ prefactor describe the same torsion field); the physical contact term $-\frac{3\kappa}{16} (J^5)^2$ is identical in both conventions.

the gravitational G used elsewhere in this paper; the subscript N is retained here to disambiguate it from the Holst-sector coefficients also entering the Lagrangian).

Step 3: Parity-Odd Effective Action.—Motivated by the Holst+non-minimal-fermion construction of Mercuri [15], in which the Nieh–Yan invariant is reconstructed and the Barbero–Immirzi parameter drops out of the classical dynamics, we introduce as a phenomenological ansatz the parity-odd term:

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

where $M = M_{\text{area-gap}} \sim M_{\text{Pl}}/\sqrt{\gamma}$ is the LQG area-gap mass scale (from the LQG area-gap $\Delta \propto \gamma \ell_P^2$, the inverse-length / mass scale is $M_\Delta \sim M_{\text{Pl}}/\sqrt{\gamma}$ up to numerical constants), α is a dimensionless coupling, and $\mathcal{F}^{IJ}[K, \mathring{R}]$ 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. (6) displays its leading contribution. (The calligraphic \mathcal{F} is reserved for this gravitational curvature; the electromagnetic field strength of Sec. IV F is written $F_{\mu\nu}$.)

In components, the leading contribution reduces to

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




FIG. 2. Energy density hierarchy from the Planck scale to the observed dark energy scale, illustrating the single-scale NDA dimensional no-go $\rho_{\text{vac}} \sim [(\alpha/M) M_{\text{Pl}}] M_{\text{Pl}}^4$ (Sec. II A 2, Appendix B). Single-scale power counting ($\Lambda \sim M_{\text{Pl}}$) pins the natural density at $\sim M_{\text{Pl}}^4$, never (meV) 4 ; no positive amplitude is derived from the ECH action, and none is needed — the $+1 \rightarrow +4$ dimensional gap is itself the no-go. The dilution waypoint quoted in the panel is the quantitative bookkeeping of Sec. XII A and Appendix B: $N_{\text{tot}} \approx 92$ with $\mathcal{D}_{\text{inf}} \approx 4 \times 10^{-122}$.

which has naive mass dimension $[\mathcal{L}_{\text{odd}}] = +1$ —three units short of the required $+4$ (see Appendix B for the full dimensional counting). We state this off-shell mismatch deliberately and do *not* claim the operator is dimension-4: with $[\alpha] = 0$, $[M] = +1$, $[e_{\mu}^J] = 0$, and $[\mathcal{F}_{IJ\rho\sigma}] = +2$, the integrand carries mass dimension $-1 + 2 = +1$, so \mathcal{L}_{odd} is genuinely dimension- $+1$ off-shell (in agreement with the referee bookkeeping that flags the mismatch). The identification $\rho_{\Lambda} = \Xi M_{\text{Pl}}^4$ is therefore governed by a *single-scale NDA dimensional no-go*: the three-unit deficit forces the missing powers to M_{Pl}^3 (the only scale in minimal ECH), so the natural density is $\sim M_{\text{Pl}}^4$, never (meV) 4 (Appendix B). This is a dimensional-impossibility argument, not a derived amplitude, and hence not circular; it is treated as such throughout.

Step 4: Parity-Odd Coefficient.—Following Freidel *et al.* [16] and Shapiro & Teixeira [20], the one-loop estimate is

$$\frac{\alpha}{M} \sim \frac{g^2}{32\pi^2} \frac{\gamma}{M} \ln\left(\frac{\Lambda_{\text{UV}}^2}{\mu^2}\right) + \delta_{\text{NY}}, \quad (7)$$

where the additive finite Nieh–Yan piece δ_{NY} carries mass dimension -1 (matching $[\alpha/M] = -1$, as required for additive consistency in Eq. (7); it is a scheme-dependent finite remainder left unestimated here), motivating the order of magnitude $[(\alpha/M) M_{\text{Pl}}] \sim 10^{-2}$. Numerically,

taking $g^2 = 4\pi\alpha_{\text{em}} \approx 0.092$ for the electromagnetic estimate, $\gamma \approx 0.274$, $M = M_{\text{Pl}}/\sqrt{\gamma}$, and the full Planck-to-TeV logarithm $\ln(\Lambda_{\text{UV}}^2/\mu^2) \approx 74$, the first term evaluates to $[(\alpha/M) M_{\text{Pl}}] \approx 3 \times 10^{-3}$ — within a factor of a few of the adopted 10^{-2} ; smaller logarithms or couplings widen the gap, which must then be carried by the unestimated δ_{NY} contribution. This is precisely why we treat α/M as a phenomenological parameter constrained by data rather than as a value derived from Eq. (7).

3. Parameter Naturalness

Required dilution of inherited rotation is naturally achieved through ~ 50 e -folds of inflation. (This ~ 50 e -fold statement concerns rotational dilution only; it is unrelated to the separate $N_{\text{tot}} \approx 92$ dark-energy dilution requirement of Sec. XII A.)

B. Black Hole Interior and Quantum Bounce

In Loop Quantum Cosmology, holonomy corrections produce a non-singular bounce at Planck-scale densities [11]:

$$H^2 = \frac{8\pi G}{3} \rho \left[1 - \frac{\rho}{\rho_{\text{crit}}} \right], \quad (8)$$

$$\rho_{\text{crit}} = \frac{3}{8\pi G \gamma^2 \Delta} = \frac{\sqrt{3}}{32\pi^2 \gamma^3} \rho_{\text{Pl}}, \quad (9)$$

with $\Delta = 4\sqrt{3}\pi\gamma\ell_P^2$ the LQG area gap. Ashtekar & Singh [11] quote the canonical LQC value $\rho_{\text{crit}} \simeq 0.41 \rho_{\text{Pl}}$ at the standard LQC area-gap choice $\gamma = 0.2375$. Substituting instead the SU(2) black-hole-entropy value $\gamma_{\text{SU}(2)} \approx 0.274$ (Eq. 2) into the same formula gives $\rho_{\text{crit}} \simeq 0.27 \rho_{\text{Pl}}$; this lower value is an internal extrapolation across counting schemes (not a value quoted in Ref. [11]), and the 0.27–0.41 ρ_{Pl} window used elsewhere in this paper should be read as a scheme-dependent range rather than as a published LQC range. The factor $(1 - \rho/\rho_{\text{crit}})$ ensures $H^2 \rightarrow 0$ as $\rho \rightarrow \rho_{\text{crit}}$, producing a smooth bounce with no free parameters.

C. Cosmic Rotation and Dark Energy

The effective cosmological constant is parameterized as:

$$\Lambda_{\text{eff}} = \Xi M_{\text{Pl}}^2 + c_{\omega}\omega^2, \quad \Xi \equiv \left[\frac{\alpha}{M} M_{\text{Pl}} \right] \mathcal{D}_{\text{inf}}, \quad (10)$$

where CMB isotropy bounds give $(\omega/H)_0 < 5 \times 10^{-11}$ [21], making rotation completely negligible. Units: Λ_{eff} carries curvature units ($[\text{mass}]^2$, the standard

Einstein-equation convention for Λ) and Ξ is dimensionless; the corresponding vacuum energy density is $\rho_\Lambda = \Lambda_{\text{eff}} M_{\text{Pl}}^2 = \Xi M_{\text{Pl}}^4$, so Ξ is the dimensionless ratio $\rho_\Lambda/M_{\text{Pl}}^4$ — consistent with the $\Xi \lesssim 10^{-123}$ identification below and the M_{Pl}^4 ansatz of Appendix B. Throughout this paper $M_{\text{Pl}} = G^{-1/2} \approx 1.22 \times 10^{19}$ GeV is the *unreduced* Planck mass; the reduced-mass distinction ($\bar{M}_{\text{Pl}} = M_{\text{Pl}}/\sqrt{8\pi}$, a factor $8\pi \approx 25$ in M_{Pl}^2) is below the order-of-magnitude resolution of every estimate in this paper. The $c_\omega \omega^2$ entry is a phenomenological bookkeeping bound with dimensionless $c_\omega = \mathcal{O}(1)$, *not* a derived isotropic vacuum term: in rotating (Bianchi-type) cosmologies vorticity sources anisotropic stress rather than an isotropic Λ , and the entry is retained solely to demonstrate negligibility, $(\omega/H)_0^2 < (5 \times 10^{-11})^2 = 2.5 \times 10^{-21}$. The dark energy scale is set by $\Xi \sim 10^{-123}$ (Sec. XII A).

We emphasize that Eq. (10) is a *phenomenological parameterization*, not a first-principles positive-amplitude derivation. The parity-odd operator (Eq. 6) has naive mass dimension +1; the identification $\rho_\Lambda = \Xi M_{\text{Pl}}^4$ is governed by a single-scale NDA dimensional no-go (Appendix B) that pins the natural density at $\sim M_{\text{Pl}}^4$, never (meV)⁴. The 13 mechanism-class constraints (Sec. IX) close all routes to obtaining an observed-scale ρ_Λ from fundamental ECH dynamics.

1. Inflationary Suppression

The contorsion dilutes as a^{-3} during inflation:

$$\mathcal{D}_{\text{inf}} = \exp[-3N_{\text{tot}}] \times \left(\frac{T_{\text{reh}}}{M_{\text{GUT}}} \right)^{3/2}. \quad (11)$$

Order-of-magnitude matching for Eq. (11).—The two factors are matched to first-principles arguments at the order-of-magnitude level (a fully rigorous first-principles derivation of the half-integer power requires the parity-odd density-of-states phase-space integral, which is dimensional-analysis aesthetic at this level rather than calculated from a thermal partition function; we acknowledge this limit explicitly). (i) The $\exp[-3N_{\text{tot}}]$ factor comes from the dilution of the torsion contribution to the effective action between bounce-time t_{bounce} and reheating-time t_{reh} . Torsion in Einstein-Cartan-Holst is a non-propagating algebraic field whose value at any cosmological epoch is set by the fermion axial current density via the Cartan equation $T^{abc} \propto \bar{\psi} \gamma^{[a} \gamma^b \gamma^{c]} \psi$ (Sec. IV A, Eq. 13). Fermion number density dilutes as a^{-3} under cosmological expansion (this is the standard cold-relic scaling for a non-relativistic species; the axial-current expectation value tracks the same dilution at the operator level, $\langle J_\mu^5 \rangle \propto n_\psi$, so $\langle J_\mu^5 \rangle$ likewise dilutes as a^{-3} at the bounce-density regime where the algebraic Cartan relation is saturated; see Hehl *et al.* [12] for the original derivation in Einstein-Cartan, and Mercuri [15] for the Holst extension). Integrating the dilution from t_{bounce} to t_{reh} gives a multiplicative fac-

tor $(a_{\text{bounce}}/a_{\text{reh}})^3 = \exp[-3N_{\text{tot}}]$ where N_{tot} is the total number of inflationary e -folds between bounce and reheating (the standard N -fold parameter). (ii) The $(T_{\text{reh}}/M_{\text{GUT}})^{3/2}$ matching coefficient connects the GUT-scale physics that fixes the parity-odd operator coefficient α/M at the bounce-density end to the reheating-temperature operator that is integrated against the late-time matter sector at the dark-energy end. The fermion number density at reheating is $n_\psi(T_{\text{reh}}) \sim T_{\text{reh}}^3$ (relativistic thermal-equilibrium limit at $T \ll M_{\text{GUT}}$), and the parity-odd operator in Eq. (10) has mass dimension +1 (Sec. II C), so the matching from M_{GUT} -scale operator-coefficient normalization to T_{reh} -scale density-of-states normalization incurs a factor of $T_{\text{reh}}/M_{\text{GUT}}$ in the operator strength and an additional $\sqrt{T_{\text{reh}}/M_{\text{GUT}}}$ from the parity-odd density-of-states factor that distinguishes the $\bar{\psi} \gamma^{[a} \gamma^b \gamma^{c]} \psi$ axial-vector contraction from the parity-even scalar contraction (the parity-odd combination carries an extra phase-space suppression at thermal equilibrium, justified here on dimensional / phase-space grounds for the axial-current variance). We emphasize that this thermal phase-space factor is *not* identifiable with the Shapiro & Teixeira [20] one-loop coefficient $\alpha_{\text{em}}/(4\pi)$ appearing in Eq. 15: the latter is a quantum loop suppression characterizing the renormalization of the Holst-sector parity-odd operator, physically *unrelated* to thermal phase-space counting. The $\sqrt{T_{\text{reh}}/M_{\text{GUT}}}$ parity-odd density-of-states factor is therefore treated as a phenomenological phase-space ansatz, not as derivable from the one-loop anomaly coefficient. The two factors compound to $(T_{\text{reh}}/M_{\text{GUT}})^{3/2}$. Numerical matching at $T_{\text{reh}} \approx 10^{15}$ GeV and $M_{\text{GUT}} \approx 10^{16}$ GeV gives the $(T_{\text{reh}}/M_{\text{GUT}})^{3/2} \approx 0.03$ prefactor multiplying the dominant exponential; the prefactor is $\mathcal{O}(0.01\text{--}0.1)$ under any T_{reh} within an order of magnitude of the GUT scale, which is the canonical inflationary-reheating regime, and therefore does not contribute to the fine-tuning hierarchy at leading order. The exponential $\exp[-3N_{\text{tot}}]$ carries the entire fine-tuning sensitivity that motivates the $\Delta N_{\text{tot}} \approx 4$ residual discussed below; the algebraic $(T_{\text{reh}}/M_{\text{GUT}})^{3/2}$ prefactor is not a tunable handle.

Matching $\rho_\Lambda \approx (2.3 \text{ meV})^4$ requires $N_{\text{tot}} \approx 92$ (a fitted parameter, not predicted); $N_{\text{tot}} \approx 92$ is the canonical value used throughout this paper, while the independent M_{Pl}^4 -to- $\rho_\Lambda^{\text{obs}}$ order-of-magnitude estimate of Appendix B gives ≈ 94 , a $\sim 2\%$ reparameterization offset that does not affect any closure conclusion. This reparameterizes the fine-tuning hierarchy from 10^{122} (the genuine $M_{\text{Pl}}^4/\rho_\Lambda^{\text{obs}}$ cosmological-constant hierarchy; see Appendix B) to $\sim 10^5$ as sensitivity to $\Delta N_{\text{tot}} \approx 4$ e -folds.

Reheating thermal-reset barrier (supporting B14).—Even granting the dilution bookkeeping above, torsion in minimal ECH is non-propagating and tracks the *instantaneous* local fermion axial current density via the Cartan algebraic equation $T^\lambda{}_{\mu\nu} \propto S^\lambda{}_{\mu\nu} \propto \langle \bar{\psi} \gamma^{[\lambda} \gamma^\mu \gamma^{\nu]} \psi \rangle$. Critically, this source is the *axial current* $\langle J_\mu^5 \rangle$, not the total fermion number density n_ψ : a thermal unpolarized fermion bath in approximate C/P -equilibrium has

$\langle J_\mu^5 \rangle \rightarrow 0$ in the mean even when $n_\psi \sim T_{\text{reh}}^3 \sim 10^{45} \text{ cm}^{-3}$ is enormous. Reheating from the inflaton drives the post-bounce plasma into precisely such a thermal regime: the chirality-flipping and depolarizing interactions that equilibrate the axial-current expectation value are *expected* to exceed the Hubble rate at $T \sim T_{\text{reh}}$ (the operative requirement is $\Gamma_{\text{wash}}(T_{\text{reh}}) > H(T_{\text{reh}})$, a condition rather than a result of the present analysis), so any *coherent* bounce-era axial-current background would be rapidly washed out toward $\langle J_\mu^5 \rangle \simeq 0$ in any regime that satisfies this inequality. The expected ordering at the GUT-scale reheating considered here ($T_{\text{reh}} \sim 10^{15} \text{ GeV}$, $H_{\text{reh}} \sim T_{\text{reh}}^2/M_{\text{Pl}} \sim 10^{11} \text{ GeV}$) involves three SM channels: (i) SM Yukawa chirality-flipping rates $\Gamma_y \sim y_f^2 T$, with the top Yukawa $y_t \sim 1$ giving $\Gamma_t/H \sim y_t^2 M_{\text{Pl}}/T \gg 1$ at $T \sim T_{\text{reh}} \sim 10^{15} \text{ GeV}$, which is the dominant rapidly-thermalizing channel at the GUT-scale reheating considered here; (ii) electroweak-sphaleron $B + L$ -violation, which is unsuppressed in the electroweak symmetric phase ($\Gamma_{\text{sph}} \sim \alpha_W^5 T$, Kuzmin–Rubakov–Shaposhnikov [22]), giving $\Gamma_{\text{sph}}/H \sim \kappa \alpha_W^5 M_{\text{Pl}}/T \gg 1$ only for $T \lesssim T_{\text{sph}} \sim 10^9\text{--}10^{10} \text{ GeV}$ (the precise crossover depends on the $\mathcal{O}(10\text{--}100)$ lattice normalization κ of the sphaleron rate [22]: the bare $\alpha_W^5 M_{\text{Pl}}/T = 1$ estimate gives $T_{\text{sph}} \sim 3 \times 10^{10} \text{ GeV}$, while the conventional $\kappa \sim 25$ lattice coefficient lowers it to $T_{\text{sph}} \sim \text{few} \times 10^9 \text{ GeV}$) — so the sphaleron channel does not itself exceed H at T_{reh} , but completes the erasure once the plasma cools through the T_{sph} regime while remaining in the symmetric phase, and is exponentially suppressed only below the electroweak phase transition; (iii) neutrino-oscillation chirality randomization, which is model-dependent (mass and mixing-spectrum dependent) and at most sub-dominant to (i)/(ii) above the EW scale. A full Boltzmann calculation of $\Gamma_{\text{wash}}(T)$ vs $H(T)$ across the bounce-to-reheating window—tracking sphaleron, top-Yukawa, and right-handed-neutrino rates simultaneously—is left to a follow-up; the present *conditional* closure statement reads: *if* $\Gamma_{\text{wash}}(T_{\text{reh}}) > H(T_{\text{reh}})$ for any one of the SM chirality-flipping channels above (the expectation, given $y_t^2 M_{\text{Pl}}/T \gg 1$ at T_{reh} , with electroweak sphalerons only exceeding H at $T \lesssim T_{\text{sph}} \sim 10^9\text{--}10^{10} \text{ GeV}$), *then* the coherent axial component is thermally reset. Because the Cartan equation is *algebraic* (no kinetic term for torsion in the minimal-ECH action), erasure of the coherent axial component is instantaneously inherited by the torsion configuration: the post-reheating mean torsion is set by the thermal expectation $\langle J_\mu^5 \rangle_T$, which vanishes in the mean for an unpolarized C/P -symmetric thermal bath; we do not assign a quantitative scale to the incoherent fluctuation residual here — the closure rests on the rate-versus-Hubble washout argument above, not on a residual-amplitude estimate. Any “memory” of bounce-era torsion stored in \mathcal{D}_{inf} is therefore overwritten by the reheating thermal reset to a coherent-mean-zero, incoherent-fluctuation-only configuration, providing a *plausible* thermodynamic erasure channel for the ECH dark-energy route that does not require the dimen-

sional bookkeeping of Appendix B or the specific value of N_{tot} ; a quantitative $\Gamma_{\text{wash}}(T)$ vs $H(T)$ computation is left to a follow-up and is not used as a primary closure here. The scenario-disambiguation is: a spin-density inflaton scenario closes via the perturbation-transparency result (B14, Sec. X); a fermion-dominated reheating scenario *additionally* closes via the conditional thermal-reset argument above. This conditional strengthening of Barrier 14 (perturbation transparency) supplies a parallel thermodynamic erasure channel for the coherent torsion-sourced dark-energy mechanism, contingent on the inequality $\Gamma_{\text{wash}} > H$ at T_{reh} being satisfied in detail. Returning to the N_{tot} bookkeeping from Sec. II C 1: we emphasize that this is bookkeeping, not progress: the residual 10^5 tracks the inverse-dilution exponential $e^{+3\Delta N_{\text{tot}}}$ ($\Delta N_{\text{tot}} \approx 4$, Sec. II C 1; i.e. $1/\mathcal{D}_{\text{inf}}$, since $\mathcal{D}_{\text{inf}} \propto e^{-3N_{\text{tot}}}$) and inherits its sensitivity from the initial-condition choice for N_{tot} (not from any reheating-rate quantity), while the fixed $[(\alpha/M) M_{\text{Pl}}] \sim 10^{-2}$ prefactor does no work on the cosmological constant problem itself. The framework has not solved the cosmological constant problem; it has only relocated the fine-tuning into inflationary initial conditions.

2. Galaxy Spin Alignment Mechanism

The parity-odd operator coupling $\alpha/M \sim 10^{-21} \text{ GeV}^{-1}$ is far too Planck-suppressed to source any observable galaxy spin asymmetry; we do not attempt an explicit mapping from this coupling to a spin-dipole amplitude, and treat the predicted effect simply as negligible relative to current sensitivity. An independent ViT-Small chirality classifier applied to the DESI Legacy DR8 galaxy population — with the dipole statistic computed on the spiral-classified high-confidence subsample ($N \approx 9.5 \times 10^5$ equivariant spirals at winning-class confidence > 0.6), not on the unselected all-galaxy sample — confirms the null at the dipole level (catalog construction, sample size, accuracy, bias-audit suite, and dipole/hemisphere/ $f_{\text{CW}}^{\text{eq}}$ significances are reported in Paper IV [23]). Galaxy spin asymmetry is not a prediction of the theory.

III. OBSERVATIONAL SIGNATURES

A. CMB E - B Cross-Correlations

The parity-odd effective action would generate CMB polarization signatures through cosmic birefringence if supplemented by a photon-sector coupling (not derived here); the quantitative benchmark is spectator-ALP phenomenology. For a spatially uniform rotation:

$$C_\ell^{EB} \approx 2\beta (C_\ell^{EE} - C_\ell^{BB}). \quad (12)$$

Eq. (12) is the small-angle, spatially uniform-rotation limit; the C_ℓ^{BB} term, dominated by gravitational lens-

ing at current sensitivities, is *not* neglected in the published β estimators whose measured values we quote below [3, 4], and this paper performs no independent *EB*-based β extraction. Connecting to a quantitative rotation angle β from the gravitational/torsion operator requires an explicit photon-torsion coupling that has not been derived here. Published measurements report $\beta_{\text{obs}} = 0.342^\circ \pm 0.094^\circ$ (WMAP+Planck [3, 4]) and $0.215^\circ \pm 0.074^\circ$ (ACT DR6 [5]); the spectator-ALP benchmark $\beta \approx 0.27^\circ\text{--}0.30^\circ$ used in this paper lies within the 1σ bands of both, and the parity-odd structure is qualitatively consistent with this observed isotropic birefringence. Spectator-ALP parameter fitting and the Na-Master pipeline validation are in companion Appendix E.

B. Galaxy Spin Asymmetry: A Confirmed Null

An independent ViT-Small chirality classifier (full bias-audit, sample size, accuracy, and dipole/hemisphere significances reported in Paper IV [23]) returns a null all-sky dipole on the spiral-classified high-confidence subsample and is in amplitude tension with Shamir’s claimed $\sim 3\%$ asymmetry by a factor of $\sim 6\text{--}12$ under that pipeline (a matched-footprint Ganalyzer-style reanalysis is required for a likelihood-level exclusion; see Paper IV [23]). The minimal ECH framework predicts a spin-dipole amplitude A_0 far below current observational sensitivity, consistent with this observed null.

MCMC verification and cosmological fits.—The $\Lambda\text{CDM}+\Delta N_{\text{eff}}$ companion analysis finds $\Delta N_{\text{eff}} \approx 0$ and recovers a Hubble constant consistent with standard ΛCDM at the Planck 2018 prior level. Sample-size and dataset details, posterior values with uncertainties, MCMC diagnostics, and the AIC/BIC model comparison are hosted entirely in the folded MCMC/reproducibility appendices (Appendix E); this theory paper carries forward only the structural conclusion (no ΔN_{eff} tension closure attributable to ECH).

IV. FOUR-ROUTE NO-GO: WHY EACH STANDARD ECH CHANNEL CLOSES

The four-route channel-level closure presented in this section is established by ruling out, in turn, the four routes by which a minimal Einstein–Cartan–Holst (ECH) sector could in principle source a parity-odd or dark-energy contribution at the level required by the observational budget of Sec. III. (This channel-level closure is logically distinct from the perturbation-transparency result of Sec. X and from the dark-energy-vs-bounce structural tension of Sec. XIV D.) We collect those four routes here and close each with standard torsion-elimination derivation for R1, standard spectator-ALP phenomenology plus the naturalness audit for R4, and ansatz-level amplitude budgets for R2–R3, constituting the no-go audit for each route. These four are the minimal-ECH

dark-energy channels most commonly discussed in the prior literature; we present them as an illustrative, explicitly *non-exhaustive* enumeration of the routes by which the minimal sector could source dark energy, not as a proven complete diffeomorphism-invariant operator basis (the omitted operators and the precise sense of ‘channel-level’ are stated in the Scope paragraph below).

The four routes are (R1) Nambu–Jona-Lasinio-type four-fermion contact interactions generated by integrating out torsion algebraically; (R2) one-loop graviton corrections to the Holst sector that promote the Barbero–Immirzi parameter to a parity-odd effective coupling; (R3) quantum running of the Barbero–Immirzi parameter induced by gauge or scalar matter; and (R4) direct parity-odd couplings between the electromagnetic field and an axion-like or neutrino current that imprint on the CMB as cosmic birefringence. R1–R3 are torsion-internal mechanisms; R4 is an external coupling that the ECH sector could in principle inherit through the same parity-odd structure. R1–R3 are closed at the amplitude level under the explicitly-labeled scaling/ansatz assumptions stated below; R4 is closed at the level of a naturalness/explanatory-deficit objection rather than an amplitude exclusion (Sec. IV F).

a. Scope: channel-level enumeration, not an operator-level basis. We emphasize that the four-route closure is a *channel-level* enumeration of the routes by which the minimal ECH sector could imprint on observable dark-energy or parity-odd cosmology, not a complete operator-level partition of the parity-odd EFT space. We flag at the outset — so that no reader mistakes a stated boundary for an undisclosed gap — that the four features a referee might reasonably challenge (operator-basis incompleteness, the +1-vs-+4 off-shell mass dimension of the dark-energy ansatz, the upper-bound EFT coefficients used to close R2–R3, and the free-coupling degeneracy of R4) are precisely the claim boundaries this paper adopts *by construction*, not defects it fails to address: the paper claims a channel-level amplitude assessment under explicitly-labeled assumptions and does not claim an operator-level theorem, a first-principles dark-energy derivation, or an amplitude no-go for R4. Each of these boundaries is stated in the abstract and re-stated at its point of use below; a reviewer seeking an operator-level basis, a derived (rather than fitted) ρ_Λ , or a rigid R4 amplitude exclusion is asking for a strictly stronger result than the one claimed here, and its absence is a scope statement, not an error. In particular, R1 (NJL parity-even four-fermion) and R4 (parity-odd ALP/axial-current CMB coupling) are not logically independent at the dimension-6 operator level: both are projections of the same torsion-elimination operator generated by the Holst-extended Einstein–Cartan action [15, 16], and two additional operators in the parity-odd sector (the Jackiw–Pi gravitational Chern–Simons term $R \wedge \tilde{R}$ and the parity-odd four-fermion partner of R1 carrying the $\gamma_{\text{BI}}/(\gamma_{\text{BI}}^2+1) \cdot 8\pi G$ coefficient) are now closed explicitly

in Sec. IV B and Sec. IV C: the four-fermion partner is the parity-odd projection of the same torsion-elimination operator as R1 and inherits R1's M_{Pl}^{-2} suppression and vanishing coherent mean field, while the gravitational Chern–Simons term is a total derivative for constant coupling (zero equation-of-motion contribution) and R4-class under any non-minimal dynamical coupling. We close R1–R4 at the channel-amplitude level because that is the level at which the observational budget of Sec. III discriminates; a full operator-level no-go requires enumerating the *complete* dimension-6 parity-odd four-fermion basis (all Fierz structures) together with the gravitational Chern–Simons invariant and a projection lemma; for the four-fermion sector this projection lemma is now provided in Appendix C, leaving only a non-minimal completion outside the established no-go. The robustness check provided by the 14-barrier closure of Sec. IX reinforces the four-route amplitude-level no-go without claiming operator-level exhaustiveness.

b. Why the finite basis is closed by one argument (minimal-ECH completeness). Although we do not claim a complete diffeomorphism-invariant partition of the full parity-odd EFT, the enumeration *is* complete for the *minimal* ECH field content by a symmetry counting that collapses to a single suppression lemma. Two structural facts, both established in the body, fix a finite basis: (F1) torsion is algebraic and non-propagating (the Cartan constraint $T^{abc} = \kappa S^{abc}$, Eq. (3)), so it sources no independent dynamical channel and returns only contact operators built from the spin current together with the metric/topological sector; and (F2) minimal fermion coupling makes the spin current totally antisymmetric ($S^{abc} \propto \varepsilon^{abcd} \bar{\psi} \gamma_d \gamma^5 \psi$), so the trace-vector and tensor torsion irreps—and any operators built from them—appear only under non-minimal couplings outside the present scope. The residual gauge- and Lorentz-invariant operators that can carry a coherent $w = -1$ density are then exhausted, at mass dimension ≤ 6 , by (i) the torsion-squared/four-fermion class (all Fierz structures VV, AA, VA, TT , sharing the single $\kappa = M_{\text{Pl}}^{-2}$ prefactor of the same torsion-elimination operator—R1 and its Holst partner are two projections, and any Fierz rearrangement is an $O(1)$ recombination that cannot alter the M_{Pl}^{-2} power), and (ii) the parity-odd curvature/topological class (Holst, Nieh–Yan, and the Jackiw–Pi gravitational Chern–Simons term), each a total derivative for constant coupling and R4-class for any dynamical coupling. Because the dimension-+1 parity-odd operator of Eq. (6) is the *least*-suppressed representative, and every higher-dimension operator carries strictly greater M_{Pl} suppression, the single-scale NDA bound of Appendix B is monotone in operator dimension and therefore *class-blind*: bounding the ceiling bounds the whole finite tower. This monotonicity is a statement about the NDA coefficients: it assumes no higher-dimension operator carries an anomalously enhanced Wilson coefficient, which holds automatically under single-scale NDA (the same single-scale assumption already flagged

in Appendix B) and is the only genericity input the argument requires. The four routes are thus not merely representative but exhaustive at the level of M_{Pl} -power-counting classes *within minimal ECH*; the only operators that evade the NDA bound are those introducing a new light scale $\mu \ll M_{\text{Pl}}$ or an exact cancellation—i.e. a non-minimal extension, which is by construction the tuning the mechanism is meant to explain (Appendix B). The fully explicit Fierz-by-Fierz projection lemma is proven in Appendix C (the generated AA and VA operators Fierz-close onto $\{SS, PP, VV, AA\}$ with all coefficients dimensionless rationals, so the M_{Pl}^{-2} power is preserved term by term); both the power-counting-class completeness and its per-operator projection are thus established, and the single residual item is the non-minimal completion. We stress that the individual dimensional mismatch (+1 vs. +4) is, taken operator by operator, an elementary power-counting statement; its significance here is not novelty of the per-operator estimate but its *basis-completeness*—the same single-scale NDA ceiling, shown monotone and class-blind, bounds the entire finite minimal-ECH operator tower at once, so no unenumerated channel can smuggle in a $(\text{meV})^4$ density without introducing a new light scale. The content is the structural impossibility across the whole basis, not the triviality of any one dimensional count.

Three technical aspects of the derivation require careful dimensional and parity accounting, addressed as follows: (a) the dimensional reconstruction of $\rho_{\Lambda}^{\text{bounce}}$ (the bounce-epoch vacuum-energy scale, defined by Eq. (B2)) in Appendix B requires an internally consistent mass-dimension accounting: the on-shell ansatz inserts bounce-curvature factors to promote $(\alpha/M) M_{\text{Pl}}^3$ (off-shell dimension +2) to $(\alpha/M) M_{\text{Pl}}^5$ (on-shell dimension +4, equivalently the grouping $[(\alpha/M) M_{\text{Pl}}] M_{\text{Pl}}^4$), whereas the local-operator-promotion route absorbs the missing dimensions into the coupling ($\alpha/M \rightarrow \alpha M_{\text{Pl}}^3/M$) off-shell (Appendix B); the choice of M_{Pl}^3 vs. M_{Pl}^5 controls the subsequent $N_{\text{tot}} \approx 92$ bookkeeping, and both readings agree at the order-of-magnitude level. The dimensional reconstruction used in the N_{tot} bookkeeping rests on the on-shell density ansatz Eq. (B2), distinct from the local-operator-promotion reading; see App. B. These two readings differ in how the missing mass-dimension is supplied: the on-shell ansatz promotes $(\alpha/M) M_{\text{Pl}}^3 \rightarrow (\alpha/M) M_{\text{Pl}}^5$ by inserting Planck-scale bounce-curvature factors on-shell, whereas the local-operator-promotion route absorbs the missing powers into the coupling coefficient ($\alpha/M \rightarrow \alpha M_{\text{Pl}}^3/M$) to restore a controlled dimension-+4 EFT operator off-shell; both are phenomenological dimensional assignments (App. B), and the $N_{\text{tot}} \approx 92$ result is common to both at the order-of-magnitude level. More generally, any alternative manifestly dimension-4 local completion of Eq. (6) must reproduce the same on-shell amplitude at the bounce scale to remain phenomenologically viable; because it is that on-shell amplitude budget—not the off-shell dimensional bookkeeping—that every barrier of Sec. IX tests, the channel-level closures

are insensitive to which off-shell completion is adopted, and the +1-vs-+4 off-shell mass-dimension status of the ansatz does not alter the scaling kinetics of the constraints. A fully symmetric dimension-4 reformulation is thus a presentational refinement of the EFT bookkeeping, not a route that evades the amplitude-budget conclusions. (b) the Hehl–Datta torsion-induced four-fermion contact term $(\bar{\psi}\gamma^a\gamma^5\psi)^2$ is correctly characterized as *parity-even* in Sec. IV A: the axial-vector current $\bar{\psi}\gamma^a\gamma^5\psi$ is a pseudovector (parity-odd component by component), but the Lorentz contraction of two such pseudovectors gives a scalar that is parity-even (each component’s parity-odd factor squared is +1); the Route 1 amplitude-suppression argument stands accordingly. (c) the Route 2 one-loop graviton-correction derivation requires that the $\Delta\theta_{\text{one-loop}}/\Delta\theta_{\text{obs}}$ ratio be expressed as a dimensionless number; the dimensionless reduction is executed in-line in Sec. IV D below (with the $H_0 \rightarrow H_0/M_{\text{Pl}}$ factor restored in the numerator), and the channel-level amplitude budget that closes Route 2 (Planck suppression by $H_0/M_{\text{Pl}} \sim 10^{-60}$ in the dimensionful form, or the equivalent dimensionless ratio after the missing factor of $1/M_{\text{Pl}}$ is restored) is unaffected at the order-of-magnitude level. The qualitative closure statement that Route 2 lies below the observed birefringence amplitude by $\gtrsim 30$ orders of magnitude survives any reasonable dimensional reconciliation.

A. Route 1 (NJL four-fermion contact): closed by Planck suppression

On the standard Einstein–Cartan side—i.e. before the Holst term is added—the Cartan algebraic equation $T^{abc} = \kappa S^{abc}$ (Eq. (3)); in the half-weight torsion convention $T^\lambda{}_{\mu\nu} = \Gamma^\lambda{}_{[\mu\nu]}$ of the original literature this reads $T^{abc} = (\kappa/2) \bar{\psi}\gamma^{[a}\gamma^b\gamma^c]\psi$ — the two map exactly, see the convention footnote at Eq. (3) — allows torsion to be integrated out exactly, generating an effective four-fermion contact term whose coefficient is the gravitational coupling $\kappa = 8\pi G$ [12, 24]. Following the standard Hehl–Datta derivation, the resulting axial–axial contact interaction is

$$\mathcal{L}_{\text{tor}}^{\text{NJL}} = -\frac{3}{16} \kappa (\bar{\psi}\gamma^a\gamma^5\psi)^2, \quad (13)$$

i.e. a four-fermion operator suppressed by M_{Pl}^{-2} and *parity-even* in the *CP*-conserving Standard Model sector. The energy density that this operator contributes at cosmologically relevant fermion densities n_ψ is bounded above by $\rho_{\text{NJL}} \sim \kappa n_\psi^2 \sim n_\psi^2/M_{\text{Pl}}^2$ (where $\kappa = 1/M_{\text{Pl}}^2$ and n_ψ has mass-dim +3, so the energy density carries the correct mass-dim +4). The closure has two independent legs. (i) *Mean-field amplitude is negligible at any cosmologically relevant baryon/electron density.* A naive order-of-magnitude estimate using dense ISM-like number densities $n_\psi \sim \mathcal{O}(10^2) \text{ cm}^{-3}$ — adopted here as a deliberately conservative high-density upper bound, not the

cosmic-mean baryon density (which is $\sim 2 \times 10^{-7} \text{ cm}^{-3}$ today and would make the bound stronger still) — converted to natural units via $\hbar c = 1.973 \times 10^{-5} \text{ eV cm}$ ($1 \text{ cm}^{-3} = (1.973 \times 10^{-5} \text{ eV})^3 \approx 7.66 \times 10^{-15} \text{ eV}^3$, so $n_\psi \approx 7.66 \times 10^{-13} \text{ eV}^3$), and with $M_{\text{Pl}} = 1.22 \times 10^{19} \text{ GeV} = 1.22 \times 10^{28} \text{ eV}$ ($M_{\text{Pl}}^2 \approx 1.49 \times 10^{56} \text{ eV}^2$), gives $\rho_{\text{NJL}} \sim n_\psi^2/M_{\text{Pl}}^2 \approx 4 \times 10^{-81} \text{ eV}^4$, i.e. roughly $1.4 \times 10^{-70} \rho_\Lambda$ for the canonical $\rho_\Lambda \approx (2.3 \text{ meV})^4 \approx 2.8 \times 10^{-11} \text{ eV}^4$ used throughout this paper — far *below* ρ_Λ , not above it. The mean-field amplitude is therefore not where any late-time dark-energy contribution could hide. (ii) *Coherent vacuum-equation-of-state structure is absent.* Even setting amplitude aside, the operator $(\bar{\psi}\gamma^a\gamma^5\psi)^2$ is *parity-even* and the parity-odd axial current satisfies $\langle J^5 \rangle \approx 0$ in an unpolarized thermal bath, so there is no coherent $w = -1$ vacuum component to source late-time acceleration. The vanishing of the mean does *not* imply that the variance $\langle J^5 J^5 \rangle$ of the four-fermion contact operator is zero — an incoherent thermal contribution from the variance is permitted — but such an incoherent contribution does not carry a coherent $w = -1$ equation of state and is in any case bounded above by the amplitude estimate of leg (i). The combined conclusion is therefore: amplitude-suppressed (by ~ 70 orders relative to ρ_Λ) and parity-even / mean-zero in the coherent dark-energy sense (the late-time dark-energy claim assessed here; thermal-era densities are addressed separately in Sec. II C 1). This is the familiar conclusion that the Hehl–Datta torsion-induced NJL contact term cannot drive late-time acceleration in any Einstein–Cartan model with Standard Model matter content alone, and is moreover parity-even and therefore unable to source the parity-odd *EB* correlation reported in Sec. III A. Adding the Holst term (see R2 below) does not relax this bound: with minimal fermion coupling the torsion-elimination map modifies the contact-term coefficient only through the bounded prefactor $\gamma^2/(\gamma^2 + 1) \in (0, 1)$ of Eq. (4), which cannot enhance the amplitude above its pure-Einstein–Cartan ($\gamma \rightarrow \infty$) value. *Closure: amplitude-suppressed and parity-even.*

B. Parity-odd four-fermion Holst partner of R1: closed by the same Planck suppression and vanishing mean field

The pure-axial contact term of Eq. (4) retains only the parity-even $J^5 J^5$ projection of the operator generated by integrating out torsion in the Holst-extended Einstein–Cartan action [15, 16]. At finite Barbero–Immirzi parameter the same torsion-elimination step also produces a parity-*odd* vector–axial cross term, the genuine “Holst partner” flagged in the Scope paragraph,

$$\mathcal{L}_{\text{int}}^{\text{VA}} \propto \frac{\gamma_{\text{BI}}}{\gamma_{\text{BI}}^2 + 1} 8\pi G J_\mu J^{5\mu}, \quad (14)$$

whose coefficient carries a single power of γ_{BI} in the numerator (parity-odd), versus the γ_{BI}^2 of the parity-

even R1 term Eq. (4). This operator does *not* open a new dark-energy route, by the same two-leg argument that closes R1, because it is literally the third projection of the identical dimension-6 torsion-elimination operator (the R1 and R4 projections are the other two, as stated in the Scope paragraph), differing only by an $O(1)$ Lorentz/parity contraction: (i) *Planck suppression is inherited verbatim*. The partner shares the $\kappa = 8\pi G = M_{\text{Pl}}^{-2}$ prefactor, so its energy density is bounded exactly as R1, $\rho_{\text{VA}} \sim \kappa \langle J \rangle \langle J^5 \rangle \lesssim n_\psi^2 / M_{\text{Pl}}^2$ — the same ~ 70 -orders-below- ρ_Λ bound of Sec. IV A; the parity-odd coefficient $\gamma_{\text{BI}} / (\gamma_{\text{BI}}^2 + 1) \leq \frac{1}{2}$ is $O(1)$ and cannot lift the amplitude. (ii) *The coherent mean field vanishes*. A coherent $w = -1$ contribution requires a nonzero vacuum expectation value; the axial current satisfies $\langle J^5 \rangle \approx 0$ in a CP -conserving, unpolarized cosmological medium (as for R1), while the vector current $\langle J \rangle$ is the net fermion-number density, which redshifts as a^{-3} ($w = 0$, not $w = -1$). The cross term $\langle J \rangle \langle J^5 \rangle$ is therefore doubly suppressed and carries no coherent $w = -1$ structure, exactly as the incoherent-variance discussion of R1 already establishes. *Closure: Planck-suppressed and mean-zero, inheriting R1's budget (Tier-III)*.

C. Jackiw–Pi gravitational Chern–Simons $R \wedge \tilde{R}$: closed as a total derivative for constant coupling; R4-class otherwise

The remaining operator named in the Scope paragraph is the Jackiw–Pi gravitational Chern–Simons term [25], $S_{\text{CS}} = \frac{1}{4} \int d^4x \vartheta *RR$ with $*RR \propto R \wedge \tilde{R}$ the Pontryagin density and ϑ an embedding/coupling field. It, too, sources no dark energy in minimal ECH: (i) *Constant coupling: a total derivative (operator-level)*. In four dimensions the Pontryagin density is exactly a total derivative, $*RR = \partial_\mu K_{\text{grav}}^\mu$ — the same identity already used at Eq. (4)'s companion discussion in Sec. X to distinguish $R \wedge \tilde{R}$ from the Holst dual (where $R \tilde{R}$ is noted to be “non-zero pointwise and a total derivative even in the presence of torsion”). For *constant* ϑ the term is therefore a pure boundary term and contributes *nothing* to the equations of motion or to ρ_Λ . This is a deductive, operator-level statement, not an amplitude estimate. (ii) *Dynamical coupling: not minimal ECH, and R4-class if adjoined*. The only way $R \wedge \tilde{R}$ can source dynamics is a non-constant ϑ carrying its own kinetic term or potential. But minimal ECH contains no dynamical pseudoscalar gravitational-Chern–Simons field: the Barbero–Immirzi parameter is a *constant* fixed by the LQG area spectrum (Barrier 7, Sec. IX), so promoting ϑ to a rolling field is a non-minimal extension outside the stated scope. If nonetheless adjoined, a ϑ with an $\sim H_0$ mass/potential tuned to yield ρ_Λ is route R4 in gravitational costume — it re-imports precisely the cosmological-constant fine-tuning that closes R4 at the naturalness/explanatory-deficit level (Sec. IV F). For the scalar-matter branch the induced parity channel (gravitational-wave / CMB bire-

fringence) is moreover the same parity channel shown to be inert by the perturbation-transparency result (Sec. X). *Closure: total derivative for constant ϑ (Tier-I, operator-level); R4-class naturalness closure for any dynamical ϑ (Tier-II), reinforced by Barrier 7 and by perturbation transparency*.

a. *Residual scope*. These two named operators are now explicitly closed, upgrading the four-route channel-level survey at exactly the two points the omitted-operator flag identified. Combined with the completeness lemma above, the enumeration is basis-complete at the level of M_{Pl} -power-counting classes *within minimal ECH*: the F1 (algebraic torsion) and F2 (totally-antisymmetric axial spin current) structural facts plus NDA monotonicity leave no power-counting class unbounded. The fully explicit *Fierz-by-Fierz* form of the projection lemma is now proven (Appendix C), completing the enumeration term by term. What remains genuinely open is only a *non-minimal* completion—a new light scale $\mu \ll M_{\text{Pl}}$ or an exact cancellation—which is by construction the tuning the mechanism is meant to explain. Genuine operator-level escape therefore requires a *non-minimal* completion—a new light scale $\mu \ll M_{\text{Pl}}$ or an exact cancellation—which is by construction the tuning the mechanism is meant to explain, and is the stated scope boundary of this no-go.

D. 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 $\gamma^{-1} e^a \wedge e^b \wedge R_{ab}$ is topological in vacuum and reduces to the Nieh–Yan density on shell once torsion is integrated out [16, 26]. 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 [15, 20]. Motivated by (but *not literally derived in*) the Holst+non-minimal-fermion construction of Mercuri and the one-loop analysis of Shapiro & Teixeira—those works establish the classical structure of the Holst term coupled to fermions, the Nieh–Yan invariant, and the one-loop running of the Holst sector, not this exact effective operator—we adopt the phenomenological one-loop parity-odd operator

$$\Gamma_{\text{one-loop}}^{\text{parity-odd}} = -\frac{1}{16\pi^2} \frac{\beta(\gamma)}{M_{\text{Pl}}} \int d^4x \sqrt{-g} \partial_\mu \vartheta_{\text{NY}}(x) J^{5\mu}, \quad (15)$$

where $\vartheta_{\text{NY}}(x)$ is the Nieh–Yan pseudoscalar (mass dimension +1; written ϑ_{NY} to distinguish it from the photon-sector spectator ALP θ of Sec. IV F; no identification between the two fields is assumed anywhere in this paper), $J^{5\mu}$ is the fermion axial current, and $\beta(\gamma)$

is a slowly varying function of γ .³ The function $\beta(\gamma)$ is written with its explicit argument throughout to distinguish this renormalization-group function from the cosmic birefringence angle β of Sec. IV F. The coefficient structure of Eq. (15) is *grounded in the explicit one-loop computation* of Shapiro & Teixeira [20], 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 $\lambda_4 = \gamma \kappa^2 (W \cdot J)$ (their Eq. 37, $\kappa^2 = 16\pi G = M_{\text{Pl}}^{-2}$), with classical coefficient $\alpha_4 = -6/(1 + \gamma^2)$ and one-loop divergence coefficients $\Omega_{44} = 81\gamma^4/[16(1 + \gamma^2)^2]$, $\Omega_{24} = 81\gamma^2/[40(1 + \gamma^2)^2]$ (their Eqs. 41–42). 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. (15). 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. $\Omega_{44}/\alpha_4 = 27\gamma^4/(32(1 + \gamma^2))$, which is $\mathcal{O}(1)$ for $\gamma \sim \mathcal{O}(1)$ and $\approx 2.5 \times 10^{-3}$ at the LQG value $\gamma \approx 0.24$], not a free normalization; and (iii) the explicit $\kappa^2 = M_{\text{Pl}}^{-2}$ on every renormalized charge, i.e. the Planck suppression in the prefactor $\beta(\gamma)/M_{\text{Pl}}$. What the one-loop analysis does *not* fix is the single *absolute* normalization: Shapiro & Teixeira show the coupled flow for $\lambda_4(t)$ and $\gamma(t)$ (their Eqs. 51, 58) is a Riccati system whose particular-solution roots are complex for real γ , so the system *has no renormalization-group fixed point* and they were “unable to solve it in a completely satisfactory way.” A fully-derived Route-2 amplitude would require a UV boundary condition plus a controlled solution of that non-perturbative flow (or a dedicated matching calculation for the exact $\partial_\mu \vartheta_{\text{NY}} J^{5\mu}$ operator). We therefore treat the *absolute* $\beta(\gamma)$ normalization as a bounded EFT input while the loop factor, Immirzi-rational coefficient, and Planck suppression are one-loop-grounded. This is strictly weaker than the Route-3 result below (whose Benedetti–Speziale β -function integrates cleanly to $|\Delta\gamma/\gamma| \approx 1.4 \times 10^{-6}$) but strictly stronger than a

free ansatz: the residual freedom is a single $\mathcal{O}(1)$ normalization, and the off-shell dimension-(+1) parity-odd operator is in any case bounded by the single-scale NDA no-go of App. B regardless of that $\mathcal{O}(1)$ coefficient, since ST’s explicit $\kappa^2 = M_{\text{Pl}}^{-2}$ confirms the missing powers are Planck powers, not a light scale. Because the closure below retains $\gtrsim 60$ orders of margin, it is insensitive to this residual $\mathcal{O}(1)$ ambiguity in $\beta(\gamma)$. The dimensionless coefficient is $\mathcal{O}(\alpha_{\text{em}}/4\pi)$ multiplied by the Planck mass to a single negative power. Explicitly, the prefactor $\beta(\gamma)/M_{\text{Pl}}$ carries mass dimension -1 (the division by M_{Pl} , not multiplication): with ϑ_{NY} of dimension $+1$, $\partial_\mu \vartheta_{\text{NY}}$ has dimension $+2$ and the axial current $J^{5\mu} = \bar{\psi} \gamma^\mu \gamma^5 \psi$ has dimension $+3$, so the integrand $[\beta(\gamma)/M_{\text{Pl}}] \partial_\mu \vartheta_{\text{NY}} J^{5\mu}$ carries dimension $-1 + 2 + 3 = +4$ and the action $\int d^4x \sqrt{-g}(\dots)$ is dimensionless, as required. Once $\partial_\mu \vartheta_{\text{NY}} \sim H \sim 10^{-33}$ eV at the present epoch is substituted, the induced birefringence accumulated between recombination and today is, in the dimensionless ratio $\Delta\theta_{\text{one-loop}}/\Delta\theta_{\text{obs}}$,

$$\begin{aligned} \frac{\Delta\theta_{\text{one-loop}}}{\Delta\theta_{\text{obs}}} &\sim \frac{\alpha_{\text{em}}}{4\pi} \frac{H_0/M_{\text{Pl}}}{M_{\text{Pl}}(\alpha/M)\beta_{\text{obs}}} \\ &\sim \frac{\alpha_{\text{em}}}{4\pi} \frac{H_0}{M_{\text{Pl}}} \frac{M}{\alpha M_{\text{Pl}}\beta_{\text{obs}}}, \end{aligned} \quad (16)$$

where $\alpha_{\text{em}}/(4\pi) \approx 5 \times 10^{-4}$ (more precisely 5.8×10^{-4} ; the order-of-magnitude closure is robust to this factor), $H_0/M_{\text{Pl}} \sim 10^{-61}$, and the R4-fitted coupling $\alpha/M \sim 10^{-21} \text{ GeV}^{-1}$ gives $M_{\text{Pl}} \cdot (\alpha/M) \sim 10^{19} \text{ GeV} \cdot 10^{-21} \text{ GeV}^{-1} = 10^{-2}$. Plugging in $\beta_{\text{obs}} = 0.342^\circ \approx 6 \times 10^{-3} \text{ rad}$ ($0.342 \times \pi/180 = 5.97 \times 10^{-3}$), the dimensionless ratio is $\Delta\theta_{\text{one-loop}}/\Delta\theta_{\text{obs}} \sim 10^{-3} \cdot 10^{-61}/(10^{-2} \cdot 6 \times 10^{-3}) \approx 10^{-60}$ (canonical evaluation of the displayed contraction; we conservatively allow up to two orders of magnitude for unmodeled higher-order loop-ordering corrections, i.e. suppression by at least 10^{-58} — an explicit conservatism allowance, *not* a derived range; the eV-vs-GeV unit conversion is exact $1 \text{ GeV} = 10^9 \text{ eV}$ and is not a source of ambiguity), i.e. the one-loop induced β is suppressed by ≈ 60 (conservatively ≥ 58) orders of magnitude relative to the observed signal. We adopt this contraction as the canonical Route-2 estimate; an alternative ordering that contracts the H_0 factor with the dimensionful coupling differently yields a deliberately loose $\sim 10^{-33}$ upper bound, not used in the closure. The canonical-bound conclusion that the one-loop induced β is amplitude-suppressed many orders of magnitude below the observed WMAP+Planck birefringence signal (with ACT DR6 follow-up) is robust to this choice; Route 2 remains exploratory framing, not load-bearing for the no-go. The Route-2 amplitude is therefore far below not only the WMAP+Planck birefringence sensitivity but the observed central value itself; the one-loop Holst-sector parity-odd term cannot account for the observed birefringence amplitude. (A naive comparison of a rotation rate $\dot{\beta}$ in eV against an angle uncertainty in eV would silently treat eV·s as dimensionless; the dimensionless reduction above avoids this and recovers the standard

³ Parity classification of Eq. (15). For ϑ_{NY} a pseudoscalar, $\partial_\mu \vartheta_{\text{NY}}$ transforms as a pseudo-co-vector; $J^{5\mu}$ is also a pseudo-vector; their Lorentz-scalar contraction $\partial_\mu \vartheta_{\text{NY}} J^{5\mu}$ is therefore parity-EVEN as a Lagrangian term. The parity-violating phenomenology in Route 2 arises not from the operator’s intrinsic P transformation but from a P-breaking *background* expectation $\langle \partial_\mu \vartheta_{\text{NY}} \rangle \neq 0$ (the time-dependent cosmological Nieh–Yan pseudoscalar selects a preferred temporal orientation, spontaneously breaking P and T). The label “parity-odd” used in the section heading and surrounding text refers to the resulting parity-violating phenomenology, not the operator’s intrinsic parity. The label is retained for consistency with the section’s established terminology. The photon-coupling chain used in the dimensionless ratio below proceeds via the standard chiral-anomaly $\partial_\mu J^{5\mu} \supset (\alpha_{\text{em}}/4\pi) F\bar{F}$ relation at the EFT level (the operator above does not itself contain an electromagnetic field strength); the resulting β estimate is treated strictly as an amplitude-budget bound, not a derived prediction.

R2 amplitude-suppression closure.) *Closure: amplitude-suppressed by M_{Pl}^{-1} and one-loop factor $\alpha_{\text{em}}/(4\pi)$.* We stress the status of this estimate explicitly, since it bears directly on the strength of the no-go: Eq. (15) is an *illustrative upper-bound amplitude budget* constructed as the natural EFT operator at the $M_{\text{Pl}}^{-1}\alpha_{\text{em}}/(4\pi)$ scale, *not* a result extracted from Mercuri [15] or Date–Kaul–Sengupta [27], which establish the classical Holst/Nieh–Yan structure but not this coefficient. Because the resulting suppression carries a margin of ~ 60 orders of magnitude, the qualitative closure is insensitive to the precise coefficient: even inflating the ansatz prefactor by $\mathcal{O}(1)$ – $\mathcal{O}(10^{10})$ (far beyond any plausible EFT enhancement) still leaves the induced β tens of orders of magnitude below the observed signal. The closure of Route 2 is therefore robust to the ansatz-level status of the operator, and we present it as an amplitude-budget bound rather than a rigorous derivation. The same logic applies to Route 3 below.

E. 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 [27]; that analysis establishes a topological interpretation of γ and motivates a γ -running in the presence of chiral asymmetry, but does *not* itself present the explicit RG equation used below. Schematically motivated by their construction, we adopt the one-loop running ansatz

$$\frac{d\gamma}{d\ln\mu} = \frac{1}{12\pi^2} (N_F^L - N_F^R) \gamma + \mathcal{O}(\gamma^2), \quad (17)$$

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. (17) only as an upper-bound EFT ansatz for the Route-3 amplitude budget and do not claim it is taken verbatim from [27]. The actual fermion-induced perturbative running of the Immirzi parameter is computed by Benedetti & Speziale [28], 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. (17) 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 [29] (their Eq. 7),

$$\mu \frac{\partial\gamma^2}{\partial\mu} = -(\gamma^2 - 1) \frac{\mu^2 \kappa^2}{(8\pi)^2} (23\gamma^2 + 5), \quad \kappa^2 = 16\pi G, \quad (18)$$

whose only fixed point is the ultraviolet-attractive $\gamma^2 = 1$ (formally outside perturbative control), with the sign of the flow set by whether $|\gamma| \gtrsim 1$ and the running driven by the radiatively-generated four-fermion interaction. The decisive feature is the explicit $\mu^2 \kappa^2 = (\mu/M_{\text{Pl}})^2$ prefactor: the running is *power-suppressed* by the renormalization scale in Planck units, not logarithmic, so the accumulated $\int \beta_{\gamma^2} d\ln\mu \propto \int \mu d\mu$ is dominated by the ultraviolet endpoint and is of order $(\mu_{\text{UV}}/M_{\text{Pl}})^2$. Integrating Eq. (18) numerically from a GUT-scale ultraviolet boundary $\mu_{\text{UV}} \sim 10^{16}$ GeV down to $\mu_{\text{IR}} \sim 1$ GeV (with γ near the LQG value $\gamma \approx 0.24$) gives, in agreement with a frozen-coefficient analytic estimate to four significant figures, $|\Delta\gamma/\gamma| \approx 1.4 \times 10^{-6}$. This is a genuine integrated running, *not* an ansatz bound: the explicit $\mu^2 \kappa^2 = (\mu/M_{\text{Pl}})^2$ prefactor makes the flow power-suppressed, so $|\Delta\gamma/\gamma| \sim (\mu_{\text{UV}}/M_{\text{Pl}})^2$ and the integral is dominated by the ultraviolet endpoint. The result is perturbatively controlled for any sub-Planckian UV boundary; only a literal Planck-scale boundary would push $|\Delta\gamma/\gamma| \rightarrow \mathcal{O}(1)$, precisely the regime Benedetti & Speziale flag as outside perturbative control (the $\gamma^2 = 1$ fixed point corresponds to a divergent four-fermion coupling), so a Planck-scale evaluation would require UV-completion input beyond the one-loop β -function. The derived physical running therefore only *strengthens* Route 3, and this GUT-UV integrated value now *replaces* the earlier ansatz-tier magnitude as the primary Route-3 estimate. Propagated to the dark-energy channel through the paper’s own $(\Delta\gamma/\gamma) \cdot (H_0/M_{\text{Pl}})$ mass-dimension suppression, the derived torsion/Immirzi contribution to ρ_Λ sits ~ 41 – 67 orders of magnitude below the observed dark-energy density $\rho_{\Lambda,\text{obs}} \approx (2.25 \text{ meV})^4$ — so the no-go closes with enormous margin, now from a *derived* integrated result rather than an ansatz upper bound. We nonetheless retain the far larger chiral-count estimate Eq. (17), $\Delta\gamma/\gamma \sim 0.3$, as a deliberately pessimistic *upper* bound in the budget below. Because the Route-3 closure below carries $\gtrsim 60$ orders of suppression margin, it is in any case insensitive to the precise coefficient. In the Standard Model, the chiral asymmetry is generated by the $SU(2)_L$ doublets. Integrating Eq. (17) over the running gives $\Delta\gamma/\gamma \approx (N_F^L - N_F^R) \ln(\mu_{\text{GUT}}/\mu_{\text{IR}})/(12\pi^2)$; with a net chiral count $(N_F^L - N_F^R) = \mathcal{O}(1)$ and a GUT-to-IR lever arm $\ln(\mu_{\text{GUT}}/\mu_{\text{IR}}) \approx 30$ – 35 ($\mu_{\text{GUT}} \sim 10^{16}$, $\mu_{\text{IR}} \sim 1$ GeV), this is numerically $\Delta\gamma/\gamma \approx 0.25$ – 0.30 (a few $\times 10^{-1}$; e.g. $32/(12\pi^2) \approx 0.27$), *not* the 10^{-2} that an earlier draft mis-stated for this expression. We therefore adopt the larger $\Delta\gamma/\gamma \sim 0.3$ as the *conservative* (least-suppressed) order-of-magnitude estimate — *not* a precisely derived value, but a conservative upper bound consistent with the $|\gamma|$ -dependent Benedetti–Speziale β -function structure recorded above (four-fermion-driven, sole fixed point at $\gamma^2 = 1$, sign set by $|\gamma| \gtrsim 1$) — and note that the Route-3 closure below is insensitive to its precise value: even this $\mathcal{O}(0.3)$ running is $\ll 1$, so γ retains its order of magnitude, and it still leaves $\gtrsim 60$ orders of suppression margin. The Holst sector amplitude

that this running can source is fixed by mass dimension: any operator built from γ , R_{ab} , e^a , and the chiral current $J^{5\mu}$ must carry dimension four, which forces a single power of M_{Pl}^{-1} in the prefactor in any cosmologically relevant scalar-curvature regime. Plugging the conservative $\Delta\gamma/\gamma \sim 0.3$ into the resulting parity-odd amplitude, the cosmologically integrated effect is suppressed by an additional factor of $(\Delta\gamma/\gamma) \cdot (H/M_{\text{Pl}}) \sim 3 \times 10^{-62}$ relative to the dimensionless parity-odd amplitude budget associated with a dark-energy-scale source, closing this route by many orders of magnitude. *Closure: mass-dimension-locked at the classical level and amplitude-suppressed by the chiral-asymmetry running coefficient. Ansatz vs derivation (R2/R3):* R3's magnitude is now a *derived integrated result*: for a sub-Planckian (GUT-scale) UV boundary the verified Benedetti–Speziale one-loop β -function Eq. (18) 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 $\kappa^2 = M_{\text{Pl}}^{-2}$ Planck suppression are fixed by the Shapiro & Teixeira [20] one-loop renormalization of the Holst+fermion sector (their Eqs. 41–42, 46); only the single absolute normalization remains a bounded EFT input, because that sector's coupled RG flow has no fixed point and is not perturbatively solvable in closed form, and the operator is bounded by the single-scale NDA no-go regardless of that $\mathcal{O}(1)$ normalization. The kept ansatz value Eq. (17) ($\Delta\gamma/\gamma \sim 0.3$) is retained only as a deliberately pessimistic upper bound; the closures of R2/R3 survive an $\mathcal{O}(1)$ inflation of the ansatz coefficient because the resulting amplitude suppression ($\sim 3 \times 10^{-62}$ vs ρ_Λ , using the pessimistic $\Delta\gamma/\gamma \sim 0.3$) leaves $\gtrsim 60$ orders of magnitude of margin against any $\mathcal{O}(1)$ rescaling of the chiral-asymmetry coefficient, and the derived 1.4×10^{-6} value only widens that margin. *Strict theoretical limitation:* the NDA one-loop operator (R2) and Immirzi-running (R3) inputs used here are bounded EFT inputs, not UV-complete flow derivations—the R2 absolute normalization requires a controlled solution to the ST Riccati system (no real fixed point; ST explicitly state they could not solve it satisfactorily), and the R3 chiral-count Eq. (17) 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 $\gtrsim 60$ -order 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.

F. Route 4 (parity-odd CMB coupling via spectator ALP or neutrino current): naturalness objection rather than amplitude no-go

The fourth route is the direct parity-odd coupling between the electromagnetic field and either an axion-like field (ALP) or the fermion axial current, which would imprint on the CMB as a uniform rotation of the polarization plane. An early cosmological-birefringence treatment of this mechanism is Lue, Wang & Kamionkowski [30]; they work with a generic pseudoscalar-photon Chern–Simons coupling $\partial_\mu\phi K^\mu$ (equivalently $\phi F\tilde{F}$ up to a total divergence), not with the specific $-\frac{1}{4}(\alpha/M)$ normalization adopted here. The operator $\mathcal{L}_{\text{CS}} \supset -\frac{1}{4}(\alpha/M)\phi\tilde{F}_{\mu\nu}F^{\mu\nu}$ (with ϕ the dim+1 canonical ALP field) is the conventional ALP–photon Chern–Simons coupling used throughout the axion-electrodynamics literature; we adopt this normalization and use [30] as an early example of its cosmological birefringence implications rather than as the source of the specific prefactor.⁴ All indices are fully contracted; the integrated-by-parts equivalent $(\alpha/M)\partial_\mu\theta K^\mu$, where $K^\mu \equiv \epsilon^{\mu\nu\rho\sigma}A_\nu F_{\rho\sigma}$ is the standard Chern–Simons 4-current and $\partial_\mu K^\mu = \frac{1}{2}\tilde{F}_{\mu\nu}F^{\mu\nu}$ recovers the parity-odd contraction, is also valid. From this operator one obtains the rotation angle

$$\beta = \frac{\alpha}{2M} \Delta\phi_{\text{rec}\rightarrow\text{today}} \sim \frac{\alpha}{2M} \sqrt{2\rho_\theta/m_\theta^2}, \quad (19)$$

where the factor $1/2$ is the standard small-rotation result for the $-\frac{1}{4}(\alpha/M)$ operator normalization (the rotation angle is half the coupling times the field excursion; derived from the helicity dispersion relation in Appendix D), ρ_θ is the energy density of the spectator field and m_θ its mass. The excursion is written in the dim+1 canonical-field form $\Delta\phi = f_a\Delta\theta$ for dimensional consistency with the operator written in ϕ -canonical form (see the single-convention footnote above the operator); the

⁴ Single-convention statement (resolving an apparent dimensional ambiguity flagged in external review). Throughout this paper there is *one* pseudoscalar field and *two* normalizations of it: the dim+1 canonical field ϕ (used in the Lagrangian operator above) and the dimensionless angle $\theta \equiv \phi/f_a$ (used in Appendix D with potential $V = m_\theta^2 f^2(1 - \cos\theta)$, Eq. (D2); $f \equiv f_a$ is the dim+1 decay constant). Dimensional accounting: $[\alpha/M] = -1$, $[\phi] = +1$, $[F\tilde{F}] = +4$, so the operator displayed above has Lagrangian-density dimension $+4$ as required. The θ -dimensionless writeup is obtained by $\phi \rightarrow f_a\theta$, which sends $-\frac{1}{4}(\alpha/M)\phi\tilde{F}F \rightarrow -\frac{1}{4}[(\alpha/M)f_a]\theta\tilde{F}F$; the bracketed quantity $(\alpha/M)f_a$ is dimensionless, exactly compensating the θ dimensionlessness. Equivalently, $\Delta\theta = \Delta\phi/f_a$ is the dimensionless excursion used in Eq. (19), and the basis-conversion to $g_{a\gamma} \equiv (\alpha_{\text{em}}c_\gamma)/(2\pi f_a)$ recorded at Eq. (D4) reproduces the identification $\alpha/M = g_{a\gamma}$. The θ -versus- ϕ alternation in the body of this paper is purely a choice of writing the same operator in either θ -dimensionless or ϕ -canonical form; no two operators are in play. This footnote is the single authoritative convention statement for the paper.

dimensionless angle $\Delta\theta$ used in Appendix D’s numerical pipeline corresponds via $\Delta\phi = f_a\Delta\theta$. The excursion estimate $\Delta\phi_{\text{rec}\rightarrow\text{today}} \sim \sqrt{2\rho_\theta}/m_\theta$ assumes a coherently displaced field whose evolution between recombination and today is monotonic, valid for $m_\theta \lesssim H_0$ (frozen or slowly rolling field); since the rotation angle depends only on the endpoint values $\theta_{\text{today}} - \theta_{\text{rec}}$, a rapidly oscillating field ($m_\theta \gg H_0$) has its present-day amplitude redshift-diluted ($\rho_\theta \propto a^{-3}$ after onset of oscillation), which *suppresses* β at fixed ρ_θ and therefore requires an even larger ρ_θ to match β_{obs} — the overshoot conclusion below is conservative with respect to the assumed $\theta(t)$ regime. Setting the present-day rotation-rate amplitude equal to the published WMAP+Planck cosmological-birefringence measurement $\beta_{\text{obs}} = 0.342^\circ \pm 0.094^\circ$ ($\sim 3.6\sigma$ from $\beta = 0$; [3, 4]; the independent ACT DR6 follow-up of Diego-Palazuelos & Komatsu [5] reports $\beta = 0.215^\circ \pm 0.074^\circ$ at $\sim 2.9\sigma$, consistent within $\sim 1.1\sigma$: $|0.342 - 0.215|/\sqrt{0.094^2 + 0.074^2} = 0.127/0.120 \approx 1.06$; these two σ values are derived under different null procedures, masks, and foreground treatments and are *not* the output of a joint analysis with a known covariance, so this $\sim 1.06\sigma$ figure is a heuristic difference assuming independent Gaussian errors, not a joint-pipeline significance) bounds α/M at $\sim 10^{-21} \text{ GeV}^{-15}$, identical to the value already quoted in Sec. II A 2; identifying the spectator field with the ECH parity-odd sector and demanding that it *also* carry the observed dark-energy density imposes a tuning constraint on the spectator-field mass m_θ that re-imports the cosmological-constant problem through the back door. From Eq. (19), $\beta = (\alpha/2M)\sqrt{2\rho_\theta}/m_\theta^2$ inverts to $\rho_\theta = 2m_\theta^2\beta^2/(\alpha/M)^2$; plugging in $\alpha/M = 10^{-21} \text{ GeV}^{-1}$, $\beta = \beta_{\text{obs}} \approx 6 \times 10^{-3} \text{ rad}$, and $m_\theta = H_0 \approx 1.5 \times 10^{-33} \text{ eV}$ gives $\rho_\theta \approx 1.6 \times 10^{-10} \text{ eV}^4 \approx 6\rho_\Lambda$ — matching the dark-energy density to within an order of magnitude — so the spectator-ALP route does *technically* reproduce the dark-energy density at the R4-fitted coupling, but only by tuning m_θ to $\sim H_0$, which is precisely the cosmological constant problem in disguise

⁵ The paper’s α/M is *not* the canonical ALP–photon Chern–Simons coupling $g_{a\gamma} \equiv (\alpha_{\text{em}} c_\gamma)/(2\pi f_a)$; the two coincide numerically at $\alpha/M = 10^{-21} \text{ GeV}^{-1}$ only after a non-trivial identification. Specifically: the paper’s $M = M_{\text{area-gap}} = M_{\text{Pl}}/\sqrt{\gamma} \approx 1.9 M_{\text{Pl}}$ (using $\gamma_{\text{SU}(2)} \approx 0.274$, Eq. (2)), and the paper’s α is the dimensionless Mercuri one-loop coefficient $\sim \alpha_{\text{em}}/(4\pi) \approx 5.8 \times 10^{-4}$ (Eq. (7)), giving the $-\frac{1}{4}$ (not $1/(2\pi)$) normalization convention. The naive identification $g_{a\gamma} = \alpha/M$ at $f_a = M_{\text{Pl}}$, $c_\gamma \sim O(1)$ yields $g_{a\gamma} \sim 10^{-22} \text{ GeV}^{-1}$, roughly $10\times$ smaller than the paper’s value. To reproduce $\alpha/M = 10^{-21} \text{ GeV}^{-1}$ in the canonical basis one therefore requires either a sub-Planckian decay constant $f_a \sim M_{\text{Pl}}/10$ or an amplified photon-coupling coefficient $c_\gamma \sim O(10)$; both are non-trivial UV-completion assumptions not derived in this paper. We adopt the Mercuri-style coupling without making either assumption and emphasize that the $10\times$ basis-conversion gap is not an internal inconsistency of R4: the R4 closure (cosmological-constant fine-tuning at $m_\theta \sim H_0$) is unaffected, and the resulting parameter α/M remains an effective phenomenological parameter constrained by data, exactly as stated at Eq. (7).

rather than its solution. R4 is therefore *not* closed by amplitude mismatch (as prior analyses claimed); it is closed by the observation that the same coupling that produces β_{obs} requires an ultralight-mass tuning $m_\theta \sim H_0$ to also produce ρ_Λ , and this tuning is the original CC fine-tuning relabelled. For any m_θ in the natural ALP range ($m_a \in [10^{-22}, 10^{-15}] \text{ eV}$) the produced $\rho_\theta \propto m_\theta^2$ overshoots ρ_Λ across the entire natural range (because the $\rho_\theta = \rho_\Lambda$ matching point lies at $m_\theta \approx 0.4 H_0 \sim H_0$ with $H_0 \approx 1.5 \times 10^{-33} \text{ eV}$, and the natural ALP range lies entirely *above* that point, so the overshoot is monotonic in m_θ and is bounded below by its lower-endpoint value (~ 22 OOM at $m_\theta \sim 10^{-22} \text{ eV}$) and grows to ~ 36 OOM at the upper endpoint $m_\theta \sim 10^{-15} \text{ eV}$); this overshoot conclusion is conditional on the one-loop estimate $\alpha/M \sim 10^{-21} \text{ GeV}^{-1}$ being rigidly bounded by the photon-Chern-Simons matching — if α/M is instead treated as a free phenomenological parameter, both β_{obs} and ρ_Λ can be matched for arbitrary m_θ by scaling $\alpha/M \propto m_\theta$ (e.g., requiring $\alpha/M \sim 10^{-10} \text{ GeV}^{-1}$ at $m_\theta \sim 10^{-22} \text{ eV}$; couplings of that size at ultralight masses are moreover in strong tension with established astrophysical ALP–photon limits from helioscope and stellar-cooling constraints, so this free-coupling direction illustrates the degeneracy structure of the matching formula rather than an open parameter direction), so the rigidity of the no-go is tied to the one-loop matching assumption rather than to ALP-mass kinematics alone: at $m_\theta \sim 10^{-22} \text{ eV}$ the overshoot is ~ 22 orders of magnitude $(m_\theta/H_0)^2 \sim (10^{11})^2 \sim 10^{22}$, and at $m_\theta \sim 10^{-15} \text{ eV}$ the overshoot is ~ 36 orders of magnitude $(m_\theta/H_0)^2 \sim (10^{18})^2 \sim 10^{36}$; the $m_\theta \sim H_0$ window where both observables are simultaneously matched has fractional width $\Delta m_\theta/m_\theta \sim 10^{-1}$, representing a dimensionful tuning of order $10^{-33} \text{ eV}/M_{\text{Pl}} \sim 10^{-61}$. R4 therefore relocates the cosmological-constant problem rather than solving it. The inequality is rigid only under the one-loop matching assumption; with α/M floated, the spectator-ALP class is recovered as a *viable parity-odd source* but is *not* a predictive dark-energy source — the model contains no first-principles explanation for why $m_\theta \sim H_0$ or for the fitted value of α/M . LiteBIRD ($\sigma(\beta) \approx 0.03^\circ$, early 2030s) [31] will tighten this bound by a factor of ~ 3 , but the bound’s structure is set by the ratio of dark-energy to birefringence amplitudes, which is dimensional and instrument-independent. *Route-4 status: a naturalness objection rather than an amplitude exclusion. A free-coupling spectator-ALP fit reproduces both β_{obs} and ρ_Λ , but minimal ECH does not derive $m_\theta \sim H_0$ or the fitted α/M ; the channel is closed at the level of an explanatory deficit, not an amplitude no-go at the operator level.*

G. Closure summary

Within the channel-level enumeration of Sec. IV (“Scope” paragraph), Routes R1–R4 cover the four parity-odd / dark-energy channels enumerated in this

paper. The two operators previously omitted from this four-channel enumeration (the Jackiw–Pi gravitational Chern–Simons term $R\tilde{R}$ and the parity-odd four-fermion partner of R1 carrying the $\gamma_{\text{BI}}/(\gamma_{\text{BI}}^2+1) \cdot 8\pi G$ coefficient) are now *also* closed, in Sec. IV B (Planck suppression + vanishing mean field, inheriting R1) and Sec. IV C (total derivative for constant coupling; R4-class otherwise); the dimension-6 parity-odd four-fermion basis is Fierz-closed and uniformly Planck-suppressed by the projection lemma of Appendix C; only a non-minimal completion lies outside the established no-go (Sec. IV Scope paragraph; abstract). Within the four enumerated channels: R1 (NJL contact) is amplitude-suppressed by M_{Pl}^{-2} and parity-even. R2 (one-loop graviton corrections) is amplitude-suppressed by M_{Pl}^{-1} and the one-loop factor $\alpha_{\text{em}}/(4\pi)$. R3 (Immirzi running) is mass-dimension-locked and additionally suppressed by the chiral-asymmetry beta function. R4 (parity-odd CMB coupling) is closed by a naturalness objection: with α/M treated as a free parameter, a spectator-ALP fit reproduces both β_{obs} and ρ_{Λ} , but minimal ECH does not derive $m_{\varphi} \sim H_0$ or the fitted α/M , so the channel closes at the level of an explanatory deficit rather than an amplitude exclusion. The NJL contact term is parametrically far below ρ_{Λ} even at dense ISM-like densities $n_{\psi} \sim 10^2 \text{ cm}^{-3}$ (~ 70 orders below ρ_{Λ} ; see Sec. IV A), parity-even with $\langle J^5 \rangle \approx 0$, and lacks any coherent $w = -1$ mean-field structure; incoherent thermal variance $\langle J^5 J^5 \rangle$ is permitted but does not source coherent dark energy. MCMC fits and the ΔN_{eff} /birefringence consistency analysis are reported (Appendix E); the present manuscript provides only the channel-level amplitude closure of the four parity-odd routes in the ECH sector.

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 III 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–R3 are Tier-III ansatz-level estimates; 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 B. 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. IX.

The phenomenological parameter α/M in Sec. II A 2 is therefore best understood as the R4-bounded coupling required to match β_{obs} , with the dark-energy density supplied by an unrelated sector (e.g. a quintessence field [32], a quintom-class two-field scenario [33], or by a positive cosmological constant), rather than by ECH-internal physics. This inversion of the prior mass-coupling lock is the core structural finding of this manuscript.

V. DATA METHODS: GALAXY SPIN ANALYSIS

Prior work.—Galaxy spin dipole analysis historically relied on published CW/CCW labels from Shamir [34, 35], who reported $\sim 1\text{--}3\%$ CW excesses. These claims have been contested [36, 37].

Our chirality classifier.—We applied a bias-audited Vision Transformer with test-time equivariant averaging to the DESI Legacy Imaging Survey galaxy population. The catalog construction, sample size, validation accuracy, bias-audit suite, equivariant CW-fraction monopole, and dipole significance are reported in Paper IV [23] and are not duplicated here. The observational conclusion is the null result of Sec. III B: the all-sky dipole is null on the spiral-classified subsample, and Shamir’s 3% claim is disfavored in amplitude by a factor of $\sim 6\text{--}12$ (matched-footprint reanalysis required for a likelihood-level exclusion; Paper IV [23]).

VI. SYSTEMATIC ANALYSIS

The galaxy spin channel is a confirmed null (Sec. III B). The CMB birefringence channel provides the surviving parity-violation evidence from the published WMAP+Planck Eskilt & Komatsu measurement $\beta_{\text{obs}} = 0.342^\circ \pm 0.094^\circ$ [4] with an independent ACT DR6 follow-up ($\beta = 0.215^\circ \pm 0.074^\circ$, Diego-Palazuelos & Komatsu [5]). For the f_{NL} channel, dominant systematic uncertainties are GR-projection effects ($\sim 20\%$ amplitude degradation at $z > 2$, Heinrich *et al.* 2024 [38] Sec. 3.4), b_{ϕ} bias-prior uncertainty ($\sigma(b_{\phi})/b_{\phi} \approx 0.2$), and photo- z marginalization, propagated through the multi-bin Fisher matrix (Sec. VII). MCMC systematics (dataset-dependent ΔN_{eff}) are in the folded MCMC/reproducibility appendices (Appendix E).

VII. FALSIFIABILITY CRITERIA

The surviving testable predictions are: (1) LiteBIRD ($\sigma(\beta) \approx 0.03^\circ$, early 2030s) will measure β to $\sigma(\beta) \approx 0.03^\circ$ and either confirm a non-zero birefringence at high significance or rule out the spectator-ALP class as the source of the WMAP+Planck birefringence signal (with ACT DR6 follow-up) (the relevant comparison is differential against the prior central value $\beta_{\text{obs}} = 0.342^\circ \pm 0.094^\circ$,

TABLE III. 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 in the paper. “Rules out” is to be read at channel-amplitude granularity (the level at which the observational budget of Sec. III discriminates), not as an operator-level theorem.

Leg	What it rules out (channel-amplitude level)	Evidentiary status
Perturbation transparency (Sec. X)	The Holst sector / Barbero–Immirzi parameter contributing to scalar or tensor perturbation observables, for canonical scalar matter	(I) Rigorous within stated scope: $T = 0$ follows from zero scalar spin density, and the Holst dual contraction vanishes by the algebraic Bianchi identity. Excludes propagating-torsion, fermion-loop, dynamical-Immirzi, non-minimal-matter sectors.
R1 (NJL four-fermion contact)	The torsion-induced contact term sourcing coherent $w = -1$ dark energy	(II)+(III) : parity-even structure ($\langle J^5 \rangle \approx 0$, no coherent mean field) is a Tier-II algebraic fact; the M_{Pl}^{-2} amplitude suppression (~ 70 orders below ρ_Λ) is a Tier-III order-of-magnitude estimate from a standard Cartan torsion-elimination derivation.
R2 (one-loop graviton corrections)	One-loop promotion of γ_{BI} reaching the observed birefringence amplitude	(III) Ansatz-level : amplitude budget suppressed by H_0/M_{Pl} and $\alpha_{\text{em}}/(4\pi)$ under an explicitly-labeled scaling ansatz; exploratory, not load-bearing.
R3 (Immirzi running)	Quantum running of γ_{BI} generating a dark-energy-relevant shift	(II)+(III) : mass-dimension lock is structural; the chiral-asymmetry beta-function bound (Benedetti–Speziale) is a Tier-III order-of-magnitude upper bound, deliberately loose and non-load-bearing ($\gtrsim 60$ orders of margin).
R4 (parity-odd CMB / spectator-ALP coupling)	Minimal ECH <i>deriving</i> the single coupling that yields both β_{obs} and ρ_Λ	(II) Structural naturalness objection , <i>not</i> an amplitude exclusion: a free- α/M ALP fit reproduces β_{obs} , but ECH supplies neither $m_\theta \sim H_0$ nor the fitted α/M , relocating the CC problem. Conditional on the on-shell scaling ansatz (App. B).

not a naive $0.27^\circ/0.03^\circ$); (2) SPHEREx (~ 2028) will test the matter-bounce prediction $f_{\text{NL}} = -35/16$ at 1.3– 2.75σ realistic significance⁶ [2, 38] via the galaxy bispectrum, simultaneously discriminating matter-bounce

from slow-roll inflation ($f_{\text{NL}} \approx 0.015$) and the Cuscuton bounce [39]; (3) MCMC parameter values (H_0 , σ_8 , ΔN_{eff}) are already consistent with standard Λ CDM, constraining the framework rather than falsifying it (details in the folded MCMC/reproducibility appendices, Appendix E).

⁶ The 1.3– 2.75σ realistic range reflects two forecast regimes: $\sigma(f_{\text{NL}}) \approx 0.7$ Fisher-ideal (raw ratio $|f_{\text{NL}}|/\sigma = 2.1875/0.7 \approx 3.13\sigma$, reduced to ~ 2.6 – 2.75σ optimistic after template-overlap correction $r \approx 0.84$ between the matter-bounce shape and the local/equilateral basis, before further GR-projection and b_ϕ degradation) and $\sigma(f_{\text{NL}}) \approx 1.0$ after GR-projection and photo- z marginalization (1.3– 2.75σ realistic). Both assume nominal SPHEREx survey volume ($f_{\text{sky}} = 0.75$, $\sim 3 \times 10^8$ galaxies). The full multi-tracer SPHEREx Fisher forecast is computed in Paper II [2], which recasts the Heinrich *et al.* $\sigma(f_{\text{NL}}^{\text{local}}) \approx 0.7$ baseline for the matter-bounce template mismatch and adopts exactly these 2.6– 2.75σ optimistic and 1.3– 2.75σ realistic (post-systematic-budget) ranges as its headline forecast; the present footnote summarizes that result rather than deriving it from the in-text $\sigma(f_{\text{NL}}) \approx 1.0$ GR-marginalized value alone (which gives the $\approx 2.2\sigma$ lower midpoint).

VIII. RELATED WORK

This work builds on rotating cosmologies (Gödel [40]), ECH theory (Hehl *et al.* [12]), Popławski’s torsion bounce and black hole universe scenario [13, 14, 41], the Holst/Nieh-Yan parity structure (Freidel *et al.* [16], Mercuri [15, 42]), and cosmic birefringence detections (Minami & Komatsu [3]). Recent independent support includes Liu *et al.* [43] (EC torsion fits the S_8 tension), Legner *et al.* [44] (torsion condensation), and Alam *et al.* [45] (non-singular bounces in modified gravity). No prior work assembles these into a single quantitative frame-

work with systematic barrier testing.

Recent developments in bounce cosmology include: Cai & Zhu [46] (GW echo signatures), Papanikolaou *et al.* [47] (PBH formation in matter bounce), and Deghani *et al.* [39] (Cuscuton bounce bispectrum).

IX. STRUCTURAL CONSTRAINTS ON DARK-ENERGY ROUTES IN MINIMAL ECH

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

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

Constraint classification.—**Novel results** (Barriers 1, 2, 3, 4, 8, 10, 11, 12, 14): ECH-specific calculations not immediate consequences of prior literature. **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 1: Mass-Coupling Lock (Foundation A)

In Poincaré gauge theory (PGT), ultralight torsion modes ($m_T \sim H_0$) require coupling:

$$g_{\text{eff}} \sim \frac{1}{M_{\text{Pl}} \sqrt{|t_3|}} \sim \frac{H_0}{M_{\text{Pl}}} \sim 10^{-61}, \quad (20)$$

where t_3 is the quadratic-torsion coupling of the PGT Lagrangian controlling the tensor-torsion mode mass, with $\sqrt{|t_3|} \sim m_T^{-1}$ for an ultralight mode ($m_T \sim H_0$), which makes g_{eff} dimensionless and yields the displayed H_0/M_{Pl} equality; the chain is a scaling ansatz of the PGT mass spectrum, labeled as such, not a derived equality. To achieve $g_{\text{eff}} \sim 1$, one needs $m_T \sim M_{\text{Pl}}$. The required fine-tuning is equivalent to the standard cosmological constant hierarchy: $\delta m_T^2/m_T^2 \sim (H_0/M_{\text{Pl}})^2 \sim 10^{-122}$.

B. Barrier 2: Topological-Shift Duality (Foundation B)

In metric-affine gravity, a duality emerges:

$$\text{Mass protection} \iff \text{No geometric fingerprint.} \quad (21)$$

Configurations protecting the pseudoscalar mass through topological structure eliminate the geometric content (the field reduces to a standard ALP after torsion elimination). Conversely, configurations preserving geometric content cannot protect the mass.

C. Barrier 3: Scalar-Tensor Universality (Foundation C)

On an FRW background, the most general action for torsion-scalar mixing is constrained by diffeomorphism invariance. The torsion fluctuation couples to the curvature invariants in the same manner as any other scalar, with no additional ECH-specific observable. Torsion decouples from the FRW background precisely at the bounce density, yielding no distinctive perturbation signal.

D. Barrier 4: Planck Suppression (Foundation D)

Disformal couplings from torsion are Planck-suppressed by factors of m_ϕ^2/M_{Pl}^2 or $(\partial\phi)^2/M_{\text{Pl}}^4$. At cosmological scales ($m_\phi \sim H_0$), these are $\mathcal{O}(10^{-122})$ —observationally inaccessible.

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

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

E. Barrier 5: Scale Separation (Foundation E)

The global vacuum integral $\int d^4x \sqrt{-g} \rho_\Lambda$ cannot be connected to the local bounce density without assuming a mechanism to store and transfer the integrated vacuum energy across ~ 92 e -folds of inflation. No such mechanism exists within minimal ECH.

F. Barrier 6: Attractor-Sensitivity Dilemma (Foundation F)

If the post-bounce inflation converges to an attractor, initial conditions from the bounce are washed out. If it is sensitive to initial conditions, inflation itself is destabilized. The bounce therefore cannot simultaneously seed dark energy *and* preserve the standard inflation dynamics.

G. Barrier 7: Parameter Immunity (Foundation G)

Cyclic vacuum selection mechanisms require γ to vary across cycles. However, γ is fixed by the LQG area spectrum at a universal value; there is no mechanism within LQG to produce a landscape of γ values from which selection could operate.

H. Barrier 8: Parity-Even Interaction (Branch H)

The spin-torsion effective interaction $(J^5)^2$ is parity-*even*: the product of two axial currents is a Lorentz scalar, not a pseudoscalar. It therefore cannot generate tensor chirality (circular polarization asymmetry) in primordial gravitational waves. This was independently

confirmed by the perturbation-transparency result (Barrier 14).

I. Barrier 9: Liouville Conservation (Branch J)

Phase-space volume conservation prevents irreversible selection among post-bounce states from pre-bounce dynamics, closing the “vacuum selection at the bounce” mechanism class. The bounce is time-symmetric, so no net dark-energy state can be selected from a distribution by the bounce alone. This barrier is a heuristic closure under explicit assumptions: closed Hamiltonian (non-dissipative) evolution through the bounce, no particle production, and no coarse-grained entropy injection. Scenarios that break these assumptions (dissipative or particle-producing bounces) evade Barrier 9 as stated and are constrained instead by the amplitude-budget arguments of Barriers 10–11; Barrier 9 is not used as a stand-alone closure of any route.

J. Barrier 10: UV→IR Specificity Dilemma (Branch L)

Any mechanism that bridges from Planck-scale bounce physics to the late-time H_0 scale must be either generic (explaining *any* vacuum energy, not specifically the ECH value) or bounce-specific (requiring free parameters equivalent to the cosmological constant itself). No mechanism achieves both simultaneously within ECH.

K. Barrier 11: Decoupling Universality (Branches L/M)

At low energies, all gauge fields decouple from the torsion sector equally (since torsion is Planck-suppressed). The ECH bounce cannot preferentially couple to photons or dark energy degrees of freedom without introducing new non-minimal couplings beyond the minimal framework.

L. Barrier 12: Vacuum Amplification Ceiling (Branch M)

Gravitational wave production from the ECH bounce is bounded above by:

$$\Omega_{\text{GW}}^{\text{ECH}}|_{\text{bounce}} \lesssim \left(\frac{\rho_{\text{crit}}}{\rho_{\text{Pl}}} \right)^2 \simeq 0.07\text{--}0.17, \quad (22)$$

where we have used the LQG-bounce critical-density window $\rho_{\text{crit}}/\rho_{\text{Pl}} \simeq 0.27\text{--}0.41$ from the Ashtekar–Singh effective-LQC status report [11]. The quadratic scaling in $\rho_{\text{crit}}/\rho_{\text{Pl}}$ is adopted here as an order-of-magnitude ceiling *ansatz* (not derived in this paper); Barrier 12 is correspondingly used only as a global ceiling, not as a precise bound. This total bounce-epoch GW energy-density fraction is not directly comparable to the present-day PTA spectral-density measurement $\Omega_{\text{GW}}(f_{\text{nHz}}) \sim 10^{-9}$, which differs by both (i) redshift / cosmological-dilution factors from the bounce epoch to today, and (ii) the integration over frequency to recover a spectral density at a given band. A quantitative comparison to NANOGrav requires propagating the bounce GW spectrum through the transfer function to the nHz band, which is deferred to a forthcoming bounce-GW dedicated paper (deferred); for the present analysis, Barrier 12 closes as a global energy-density-fraction ceiling rather than a direct NANOGrav exclusion.

M. Barrier 13: Gravitational Democracy (Branches N/O)

Torsion couples democratically to all spin-1/2 matter species. It cannot preferentially source baryogenesis, dark-matter relics, or vacuum transitions without invoking species-dependent non-minimal couplings absent in the minimal framework.

N. Barrier 14: Perturbation Transparency

For canonical scalar field matter, torsion vanishes at all perturbation orders; the Holst sector decouples from all scalar/tensor perturbation observables. This is elaborated in Sec. X: the scalar-sector proof is in Sec. XB and the explicit Holst-term verification at all perturbation orders is in Sec. XD.

X. THE PERTURBATION-TRANSPARENCY RESULT

A. Statement

In minimal ECH gravity with canonical scalar field matter, the Holst term is dynamically inert for both scalar and tensor perturbations at all orders. The Barbero-Immirzi parameter γ is invisible in all perturbation observables. This generalizes Hehl *et al.* (1976) [12] to the Holst sector and to all perturbation orders. The restriction to the torsion-free ($T = 0$) branch is not an additional assumption but a *consequence* of the matter content: canonical scalar matter carries zero spin density (Step 1 below), hence in Einstein–Cartan theory sources no torsion (Step 2), and the connection reduces to Levi-Civita identically at all perturbation orders. We emphasize that the “all orders” statement does *not* rest on an order-by-order component expansion of the linearized Holst/Cartan action: it follows because Step 2 is an *algebraic* (non-derivative) Cartan constraint, so $T = 0$ holds exactly and perturbatively unmodified, and Step 4 is a pointwise *identity* (the algebraic Bianchi identity $R_{\mu[\nu\rho\sigma]} = 0$, which holds for any torsion-free connection at every field configuration). Because both load-bearing steps are exact identities rather than truncated expansions, no separate order-by-order verification is required; we display the leading (second-order) expansion in Sec. XD only as an explicit check of the identity, not as the origin of the all-orders claim. The propagating-torsion, dynamical-Immirzi-field, fermion-loop, and non-minimal-matter sectors — where $T \neq 0$ and the result need not hold — are explicitly outside the stated scope; the theorem is a statement about the scalar-matter sector of ECH, which is precisely the Levi-Civita branch, so the $T = 0$ hypothesis is the theorem’s domain rather than a gap in it.

B. Proof (Scalar Sector)

1. **Zero spin density.** A canonical scalar field has zero spin density.
2. **Zero torsion.** In Einstein-Cartan theory, $T^\lambda{}_{\mu\nu} = 8\pi G S^\lambda{}_{\mu\nu} + \dots$. With $S = 0$, $T^\lambda{}_{\mu\nu} = 0$ at all perturbation orders.
3. **Connection reduces to Levi-Civita.** $\Gamma^\lambda{}_{\mu\nu} = \overset{\circ}{\Gamma}{}^\lambda{}_{\mu\nu}$.
4. **Holst term vanishes by the first Bianchi identity.** The Holst term evaluated with the Levi-Civita connection gives $\frac{1}{2}\epsilon^{\mu\nu\rho\sigma} R_{\mu\nu\rho\sigma}(\overset{\circ}{\Gamma})$, which on a torsion-free connection *vanishes identically* by the first (algebraic) Bianchi identity $R_{\mu[\nu\rho\sigma]} = 0$: the cyclic-sum identity $R_{\mu\nu\rho\sigma} + R_{\mu\rho\sigma\nu} + R_{\mu\sigma\nu\rho} = 0$ contracted with the totally antisymmetric $\epsilon^{\mu\nu\rho\sigma}$ leaves no non-trivial component. The algebraic

Bianchi identity $R_{\mu[\nu\rho\sigma]} = 0$ holds for any torsionless connection, independently of metric compatibility; non-metricity does not invalidate the identity provided $T = 0$. The Holst dual contraction is therefore identically zero on the Levi-Civita connection (not merely a boundary term), so it contributes nothing to the action at any order. This Bianchi-vanishing is distinct from the Pontryagin density $\propto R\tilde{R}$ (a two-curvature topological invariant) — see the explicit verification below.

5. No equations of motion. A total derivative contributes nothing to variational equations at all orders. (This step is logically distinct from the pointwise Bianchi-vanishing of the previous step: at $T = 0$ the stronger pointwise vanishing already applies; the total-derivative statement covers the residual Nieh–Yan boundary term $d(e_I \wedge T^I)$ at nonzero torsion.)

C. Extension to Tensor Sector

The same five steps apply to tensor perturbations. With $T = 0$, the tensor perturbation equation:

$$h''_{ij} + 2\mathcal{H}h'_{ij} + k^2 h_{ij} = 0 \quad (23)$$

$$\mathcal{R}_H(\tilde{\Gamma}) \equiv \frac{1}{2}\epsilon^{\mu\nu\rho\sigma} R_{\mu\nu\rho\sigma}(\tilde{\Gamma}) = 0 \quad (\text{identically, by the first (algebraic) Bianchi identity}). \quad (25)$$

This is the Bianchi-vanishing of the Holst dual contraction on a torsion-free connection: in differential-form language $e^I \wedge e^J \wedge R_{IJ} = -\text{NY} + T^I \wedge T_I$, where $\text{NY} \equiv d(e_I \wedge T^I)$ is the exact Nieh–Yan density (a total derivative of the torsion boundary term), and both pieces vanish at $T = 0$ *pointwise* — not merely up to a boundary term: $\text{NY}|_{T=0} = d(0) = 0$ and $T^I \wedge T_I|_{T=0} = 0$, with the Bianchi cancellation above being the operative reason the Holst dual vanishes pointwise in the torsionless sector.⁷ *This must be carefully distinguished from the Pontryagin density* $\frac{1}{4}\epsilon^{\mu\nu\rho\sigma} R_{\mu\nu}{}^{\alpha\beta} R_{\rho\sigma\alpha\beta} \propto R\tilde{R}$, which involves *two* curvature tensors and is a separate true topological invariant — non-zero pointwise and a total derivative even in the presence of torsion. The Holst dual contraction has only *one* curvature and is therefore not the Pontryagin density; on $T = 0$ it is identically zero by Bianchi,

⁷ Explicit decomposition: the Holst integrand satisfies $e^I \wedge e^J \wedge R_{IJ}(\Gamma) = -d(e_I \wedge T^I) + T^I \wedge T_I$, i.e. $e \wedge e \wedge R = -\text{NY} + T \wedge T$, where $\text{NY} \equiv d(e_I \wedge T^I)$ is the Nieh–Yan boundary form and $T^I = de^I + \omega^I{}_J \wedge e^J$ is the torsion two-form. At $T = 0$ (canonical scalar matter, torsion-free branch), both terms vanish *pointwise*: $\text{NY}|_{T=0} = d(0) = 0$ and $T^I \wedge T_I|_{T=0} = 0$. The Holst dual contraction is therefore identically zero on the Levi-Civita connection, not merely a total derivative.

(primes denote derivatives with respect to conformal time η , and $\mathcal{H} \equiv a'/a$ is the conformal Hubble rate; Fourier modes $h_{ij}(\mathbf{k}, \eta)$ follow the standard $e^{i\mathbf{k}\cdot\mathbf{x}}$ convention for the transverse-traceless amplitude; in cosmic time, using $dt = a d\eta$, the equivalent form is $\ddot{h}_{ij} + 3H\dot{h}_{ij} + (k^2/a^2)h_{ij} = 0$) has no parity-dependent modifications. Left and right circular polarization modes propagate identically:

$$v_R(k, \eta) = v_L(k, \eta) \quad \Rightarrow \quad \Delta v = 0 \quad (\text{identically}). \quad (24)$$

No GW birefringence, no tensor chirality, no TB/EB CMB parity violation from the ECH mechanism.

D. Explicit Verification: The Holst Term in Perturbation Theory

Expanding to second order, the Holst dual evaluates on the Levi-Civita connection ($T = 0$) as:

contributing nothing to the variational equations of motion at any perturbation order. In particular, the cubic action for ζ (which determines the bispectrum) receives zero contribution from the Holst term. The bispectrum is therefore identical to the standard GR result.

E. What Would Break the Transparency

The transparency result fails if: (1) matter includes fermions with nonzero spin density; (2) the gravitational action includes kinetic terms for torsion (Poincaré gauge theory); (3) non-minimal derivative couplings between the Holst sector and matter are introduced.

Quantum scope limitation.—The perturbation-transparency result is established at the *classical* level: it follows from the algebraic Bianchi identity $R_{\mu[\nu\rho\sigma]} = 0$ on the torsion-free Levi-Civita connection, which is an exact classical identity, and from the algebraic (non-derivative) Cartan constraint that sources no torsion for canonical scalar matter. Whether classical Holst-sector decoupling survives quantum loop corrections—in particular, whether graviton loops, matter loops, or chiral/gravitational anomalies could reintroduce propagating torsion degrees of freedom or source an effective

Holst contribution at the quantum level—is *not* claimed here and lies outside the stated scope of this result. The perturbation-transparency theorem is a statement about the classical scalar-matter sector of ECH; its quantum extension would require a separate analysis of the one-loop effective action in the Holst+scalar sector.

F. Implications

The perturbation-transparency result establishes a clean dichotomy:

- *Perturbation observables* (C_ℓ^{TT} , C_ℓ^{EE} , P_k , bispectrum): Identical to standard GR. No ECH modifications at any order.
- *Nonperturbative parity channels* (ALP birefringence, primordial GW chirality): Parity-sensitive channels (model-dependent tests of γ_{BI} only under a derived γ_{BI} -dependent photon or tensor-parity coupling).

G. Discrimination Among Bouncing Cosmologies

The matter bounce prediction $f_{\text{NL}} = -35/16$ is the strongest discriminator: it provides $\sigma(f_{\text{NL}}) \approx 0.7$ (*ideal-survey* Fisher) to $\sigma(f_{\text{NL}}) \approx 1.0$ (*degraded with GR-projection + photo-z systematics*) from SPHEREx, yielding 1.3–2.75 σ model separation. NANOGrav model comparison: $\gamma_{\text{PTA}} = 2.567 \pm 0.382$ from real-KDE reanalysis of the 15-yr free-spectrum data (GPU MCMC, companion [48]; here γ_{PTA} is the GWB power-law spectral index, distinct from the Barbero-Immirzi parameter γ defined in Eq. 1). The matter-bounce prediction $\gamma_{\text{PTA}} = 3.0$ sits at +1.13 σ above the posterior mean, consistent with the data within standard frequentist tolerance. The current real-KDE GPU MCMC gives $\gamma_{\text{PTA}} = 2.567 \pm 0.382$ (Paper III § 6 [48]).

XI. THE HYBRID DARK-ENERGY LOOPHOLE

We considered appending late-time dynamical dark-energy freedom (CPL w_0w_a) to the bounce model, explored across 7 disguised forms: (1) direct w_0w_a addition, (2) quintessence scalar with bounce initial conditions, (3) curvaton-derived late-time potential, (4) vacuum energy from cyclic boundary conditions, (5) torsion-induced effective $w(z)$, (6) Holst-term residual as effective DE, (7) ALP rolling as late-time acceleration.

All 7 forms were assessed at the theoretical level only: adding w_0w_a to a bounce model produces (at the level of the theoretical fit-improvement calculation) the same fit-improvement structure as adding w_0w_a to Λ CDM, with no additional theoretical content from the bounce. We emphasize that this is a theoretical-structure conclusion, *not* a quantitative posterior-preference rejection: the dedicated DESI DR2 + Planck NPIPE + Pantheon+

+ DES-SN5YR cobaya chain with the free w_0w_a extension required to quantitatively test these 7 forms is described in the folded MCMC/reproducibility appendices (Appendix E) Table IV row “DESI DR2 w_0w_a (new)” and has not yet converged to the standard publication-quality target $\hat{R} - 1 < 10^{-2}$ (see Appendix E Sec. V for the chain status). The present program’s MCMC analysis uses stock CAMB with ΔN_{eff} only and hosts zero free- w_0w_a samples; the qualitative theoretical conclusion that ECH cannot generate w_0w_a as an additional output beyond what Λ CDM + w_0w_a already accommodates therefore stands as a *theoretical structural observation* rather than as a posterior-preference quantitative rejection. A quantitative rejection (or confirmation) is gated on convergence of the dedicated folded-appendix chain (Appendix E).

XII. DISCUSSION

A. The Inflationary Suppression Factor

The constant contribution to Λ_{eff} emerges from the interplay between the parity-odd spin-torsion interaction and inflationary dilution:

$$\Xi \equiv \left[\frac{\alpha}{M} M_{\text{Pl}} \right] \times \mathcal{D}_{\text{inf}}, \quad (26)$$

a dimensionless quantity of order the observed hierarchy $\rho_\Lambda^{\text{obs}}/M_{\text{Pl}}^4 \sim 10^{-122}$, decomposed as $[(\alpha/M)M_{\text{Pl}}] \times \mathcal{D}_{\text{inf}} \sim 10^{-2} \times \mathcal{D}_{\text{inf}}$; with the fitted value $N_{\text{tot}} \approx 92$ (Sec. II C 1), Eq. (11) gives $\mathcal{D}_{\text{inf}} = e^{-3N_{\text{tot}}(T_{\text{reh}}/M_{\text{GUT}})^{3/2}} \approx e^{-276} \times 0.03 \approx 4 \times 10^{-122}$, so that $\Xi \sim 4 \times 10^{-124}$ to order of magnitude. We stress that N_{tot} is a fitted parameter (Sec. II C 1), tuned so that Ξ reproduces $\rho_\Lambda^{\text{obs}}/M_{\text{Pl}}^4$; the order-of-magnitude residual between the raw ansatz product and the target hierarchy is exactly the $\Delta N_{\text{tot}} \approx 4$ sensitivity discussed there, not an independent prediction.

Physical-versus-mathematical scope of \mathcal{D}_{inf} .—While N_{tot} controls the mathematical $e^{-3N_{\text{tot}}}$ ansatz used in the above bookkeeping, the *physical* reheating thermal-reset barrier (supporting B14; see Sec. II C 1, “Reheating thermal-reset barrier” paragraph) already closes the bounce-era-memory dilution channel: non-propagating torsion is sourced algebraically by the instantaneous axial-current expectation value, and the post-reheating coherent axial component is washed to zero by chirality-flipping and depolarizing thermal interactions that equilibrate the axial-current expectation value. The \mathcal{D}_{inf} exponential is therefore mathematical scaffolding for an order-of-magnitude parameterization of a hypothetical un-reset channel rather than a physically operative dilution mechanism; the entries in this section are retained as parameterization-of-fine-tuning diagnostics, not as a viable dynamical channel.

The “fine-tuning reduction from 10^{122} to 10^5 ” is a reparameterization as sensitivity to N_{tot} (the total number of

TABLE V. Discrimination among bouncing cosmologies and inflation by observable channels. The symbol \checkmark (check) denotes the mechanism produces the prediction; \times denotes it does not; “—” denotes not applicable or not computed.

Model	$f_{\text{NL}} = -35/16$	ALP birefringence	γ_{PTA} (real-KDE)	$w_0 w_a$ DESI
Matter bounce (any host; not ECH-specific)	\checkmark	(spectator)	\checkmark	not tested [‡]
Slow-roll inflation	\times ($f_{\text{NL}} \approx 0.015$)	(spectator)	\times	not tested [‡]
Quintom-B	\times	(spectator)	—	consistent [†]
Cuscuton bounce	\times ($f_{\text{NL}} \approx 0$)	(spectator)	—	not tested [‡]
Ekpyrotic	\times ($f_{\text{NL}} \sim -5$)	(spectator)	—	not tested [‡]

[†]Quintom-B can in principle accommodate the DESI $w_0 w_a$ evidence; the MCMC analysis in the folded MCMC/reproducibility appendices (Appendix E) was not extended to the $w_0 w_a$ parameter space, so this row is reported as “consistent at the model level” rather than a posterior-preference \checkmark .

[‡]A free- $w_0 w_a$ posterior analysis was not completed for the present paper; no posterior-preference claim is made for or against the DESI $w_0 w_a$ evidence on the basis of the present table. The frozen MCMC posteriors hosted in the folded MCMC/reproducibility appendices cover three $\Lambda\text{CDM} + \Delta N_{\text{eff}}$ dataset combinations only, and the asymmetry between the Quintom-B accommodation row and the others is one of theoretical accommodation, not of fit quality measured in this program.

inflationary e -folds), not a resolution of the cosmological constant problem. The exponential $\mathcal{D}_{\text{inf}} \propto e^{-3N_{\text{tot}}}$ structure of Eq. (11) makes N_{tot} the single controlling parameter analytically: the residual 10^5 tracks $e^{+3\Delta N_{\text{tot}}}$ for $\Delta N_{\text{tot}} \approx 4$ e -folds (the dilution ansatz $\mathcal{D}_{\text{inf}} \propto e^{-3N_{\text{tot}}}$ implies the fine-tuning score $\propto 1/\mathcal{D}_{\text{inf}} \propto e^{+3N_{\text{tot}}}$, so the *score* rescales as $e^{+3\Delta N_{\text{tot}}}$ when N_{tot} is shifted by ΔN_{tot}), while the order-unity prefactors enter at most logarithmically. We emphasize that the $(T_{\text{reh}}/M_{\text{GUT}})^{3/2}$ prefactor itself is matched at the order-of-magnitude level rather than calculated from a thermal partition function (Sec. IIC 1); the 10^5 residual therefore inherits the same order-of-magnitude status, and the “reduction from 10^{122} to 10^5 ” should be read as a qualitative dimensional rearrangement rather than a quantitative bookkeeping result.

Caveat on the $(T_{\text{reh}}/M_{\text{GUT}})^{3/2}$ prefactor.—The $(T_{\text{reh}}/M_{\text{GUT}})^{3/2}$ prefactor used here is the dimensional-analysis-aesthetic estimate from naive scaling of the matter-bounce effective theory; a first-principles derivation requires the full bounce-junction matching that lies outside the scope of this paper. The $N_{\text{tot}} \approx 92$ e -fold structural-tension result therefore carries an order-of-magnitude uncertainty inherited from this prefactor. The sign and qualitative conclusion—that the matter-bounce mechanism cannot generate dark energy without re-importing cosmological-constant fine-tuning—survives the OOM uncertainty, because the surplus required to close the gap is ~ 14 e -folds (the fine-tuning-gap surplus of the present order-of-magnitude argument; this is a different quantity from — and smaller than — the $N_{\text{tot}} - N_{\text{exit}} \approx 32$ e -fold signal-erasure differential of Sec. XIV D, so the conclusion holds a fortiori for the larger differential), whereas the quasi-dust ε -correction-driven prefactor adjustment (the $\mathcal{O}(\varepsilon - 3/2)$ equation-of-state correction of the matter-bounce contraction, $\varepsilon = 3(1 + w)/2$, defined in the companion f_{NL} forecast paper) is $\lesssim 1$ e -fold. A rigorous derivation of the prefactor from the bounce-junction matching is deferred to future work and would not change the structural-tension verdict.

B. Theoretical Implications

Four routes to deriving ρ_Λ with $w = -1$ from first principles were tested: (i) NJL condensate, (ii) one-loop fermion effective action, (iii) dynamical Immirzi field, (iv) parity-sensitive CMB phenomenology. R1 closes via the standard published derivation; R2–R3 close at the amplitude level under explicitly-labeled scaling/ansatz assumptions; R4 closes at the naturalness level. The condensate route is closed at the amplitude level (Sec. IV A): the NJL contact term is Planck-suppressed ($\rho_{\text{NJL}} \sim n_\psi^2/M_{\text{Pl}}^2 \approx 4 \times 10^{-81} \text{eV}^4$, ~ 70 orders below ρ_Λ) and parity-even. The one-loop route is amplitude-closed under the explicitly-labeled EFT scaling ansatz. The dynamical Immirzi field reduces to a standard axion-like particle with $w = +1$ (stiff matter). The parity assessment finds no photon coupling in the minimal framework: the parity-odd one-loop operator of Sec. IV D couples the Nieh–Yan pseudoscalar to the fermion axial current, $\partial_\mu \vartheta_{\text{NY}} J^{5\mu}$, with no $F\tilde{F}$ term, so a CMB-birefringence amplitude only arises through the model-dependent $\partial_\mu J^{5\mu} \supset (\alpha_{\text{em}}/4\pi)F\tilde{F}$ chiral-anomaly chain, treated as an amplitude-budget bound rather than a derived prediction (Sec. IV D footnote).

Spectator-ALP birefringence.—A spectator ALP with $f_a \sim M_{\text{Pl}}$, $m \sim H_0$ is consistent with the published WMAP+Planck cosmological-birefringence signal ($\beta = 0.342^\circ \pm 0.094^\circ$, $\sim 3.6\sigma$ from $\beta = 0$, Eskilt & Komatsu [4]; an independent ACT DR6 follow-up [5] reports $\beta = 0.215^\circ \pm 0.074^\circ$ at $\sim 2.9\sigma$) at $f_a \sim M_{\text{Pl}}$ and $\theta_i \sim \mathcal{O}(1)$ without additional ALP-naturalness fine-tuning beyond the $m_\theta \sim H_0$ ultralight-mass tuning admitted in §IV F (a cosmological-constant-class tuning rather than an ALP-specific one — the same point is also surfaced in the structural-tension discussion of §XIV D and §XI). The same $\beta \approx 0.27^\circ$ prediction arises in standard GR with an identical ALP; it is not a distinctive ECH prediction. Full ALP MCMC parameter fitting (9,720 accepted samples, $\hat{R} - 1 < 0.01$) and LiteBIRD forecast are in the folded

MCMC/reproducibility appendices (Appendix E).

XIII. SURVIVING ECH-INDEPENDENT CLASS TESTS

Despite the channel-level closure of the four enumerated minimal-ECH dark-energy routes (13 mechanism-class constraints under the stated assumptions), the broader bounce-cosmology program retains two fully testable ECH-independent class-level predictions:

(1) **Matter-bounce** $f_{\text{NL}} = -35/16$.—The matter-dominated contracting phase of a *scalar-only* $w = 0$ matter-bounce (the bounce-class observable, not specific to ECH) produces a minimally-parameterized local-type non-Gaussianity $f_{\text{NL}} = -35/16$ [1]. We adopt $-35/16$ as the corrected central value: the historical Cai *et al.* $-35/8$ [1] is traced in Paper II [2] to a spurious $+(99/128)\sum_i k_i^3$ term in their final combined polynomial and is superseded by the $-35/16$ resolution (in agreement with the independent general- c_s derivation of Li *et al.*). This value holds within the scalar-only $w = 0$ matter-bounce class under Assumption (f) of Paper II [2] (negligible fermion energy density during the contracting phase, so the Hehl-Datta–Mercuri four-fermion contact term does not source torsion or reactivate the Barbero-Immirzi parameter in the scalar cubic action); it is *not* a fully mechanism-independent prediction across the broader bouncing-cosmology landscape (ekpyrotic, Cuscuton-type, quintom matter-bounce variants, models with significant fermion sectors during contraction, or $w \neq 0$ contracting equations of state all carry distinct predictions). Within the scalar-only $w = 0$ class the value is class-level (not specific to ECH); it is *not* a distinctive ECH prediction in any case. The full multi-bin Fisher forecast, SPHEREx parameter sensitivity (bispectrum-only $\sigma(f_{\text{NL}}) \approx 0.7$ from Heinrich *et al.* 2024, leading to 1.3–2.75 σ post-systematic-budget significance), and anomaly-optimized multi-tracer strategy are in Paper II [2]; the present paper does not perform an independent SPHEREx Fisher computation and the 1.3–2.75 σ figure is reported here only as a cross-reference. SPHEREx first science data ~ 2028 .

(2) **Spectator-ALP birefringence** $\beta \approx 0.27^\circ$.—We present this as a **consistency check**, not as a prediction: the parity-odd coefficient α/M is fitted, not derived from first principles, and an ALP with $f_a \sim M_{\text{Pl}}$, $m \sim H_0$ chosen to land near the observed signal is by construction inside its 1σ band. The *quantitative* prediction is the signature structure (achromatic uniform rotation, the *EB/TB* pattern, and consistency across frequencies and experiments) rather than the central value of β . The same ALP setup arises identically in standard GR with the same parameters, so this is not a distinctive ECH prediction. LiteBIRD ($\sigma(\beta) \approx 0.03^\circ$, early 2030s) will either confirm a non-zero β at high significance or exclude this uniform spectator-ALP benchmark ($f_a \sim M_{\text{Pl}}$, $m \sim H_0$) as the explanation of the current WMAP+Planck cen-

tral value; either outcome is informative independent of ECH.

Structural incompatibility.—An open constraint is the incompatibility between the dark-energy suppression mechanism ($N_{\text{tot}} \approx 92$ e -folds required) and the $f_{\text{NL}} = -35/16$ prediction (*definitively* erased once $N_{\text{tot}} - N_{\text{exit}} \gtrsim N_{\text{coh}} \sim \mathcal{O}(\text{few})$ at SPHEREx-relevant scales; the SPHEREx accessible wavenumbers $k \sim 10^{-4} - 10^{-1} h/\text{Mpc}$ map to bounce-era *physical* scales $k_{\text{bounce}}^{\text{phys}} \sim k_{\text{SPHEREx}} e^{N_{\text{tot}} - N_{\text{exit}}} \sim e^{32} k_{\text{SPHEREx}}$ at $N_{\text{tot}} \sim 92$ and $N_{\text{exit}} \sim 60$ (the relative e -fold differential between bounce and CMB horizon-exit; comoving wavenumbers k are constant by definition and only physical scales scale with $a^{-1} \propto e^{-N}$), deep inside the inflationary subhorizon regime where the surviving bispectrum signal is purely vacuum-inflationary, not matter-bounce contraction-mode). This is detailed in Sec. XIV D. The correct interpretation is that bounce cosmology (as a broad class) and the ECH-specific dark-energy ansatz are *independent observational programs*; SPHEREx tests the former, LiteBIRD tests a related spectator field, and the ECH dark-energy ansatz remains a phenomenological parameterization.

XIV. LIMITATIONS AND FUTURE DIRECTIONS

A. Current Limitations

1. Theoretical

- *Phenomenological α/M* : Not derived from first principles; the one-loop estimate motivates existence and order of magnitude, but the finite part depends on the γ_5 regularization scheme (Sec. II A 2, [20]).
- *Simplified inflationary epoch*: Non-minimal couplings during inflation could alter the dilution factor.
- *Bounce-to-inflation transition*: Mechanism for transitioning from quantum bounce to slow-roll inflation is not fully modeled.

2. Observational

- *Galaxy spin*: Null confirmed at the dipole level (Paper IV [23]), consistent with the > 100 -orders-of-magnitude underprediction by the ECH coupling.
- *MCMC proxy*: Stock CAMB with ΔN_{eff} is a phenomenological proxy, not a bespoke spin-torsion Boltzmann module. Full MCMC details and convergence diagnostics are in the folded MCMC/reproducibility appendices (Appendix E).

B. Robustness to Galaxy Spin Null Results

The galaxy spin channel is a confirmed null at the dipole level (full quantitative chirality results in Paper IV [23]), *consistent* with the framework (which underpredicts A_0 by > 100 orders of magnitude). The parity-violation case rests entirely on CMB birefringence from the published WMAP+Planck Eskilt & Komatsu measurement [4] with an independent ACT DR6 follow-up (Diego-Palazuelos & Komatsu [5]).

C. Discriminating Observational Channels

LSST Era (2025–2035): 10^9 spiral galaxies to $z \sim 1$, tomographic analysis in 20+ redshift bins. *CMB Experiments*: LiteBIRD, CMB-S4, and future concepts (PICO, CMB-HD). *Bounce cosmology beyond ECH*: Papanikolaou *et al.* [47] showed that asymmetric matter bounces can produce asteroid-mass PBHs as dark matter candidates with induced GWs detectable by LISA and Einstein Telescope.

D. Structural Tension: Dark Energy vs. Bounce

f_{NL}

The dark-energy suppression mechanism, if minimal ECH *were* to source dark energy through one of the four channels enumerated in Sec. IV, would require $N_{\text{tot}} \approx 92$ post-bounce e -folds; the matter-bounce $f_{\text{NL}} = -35/16$ would be *definitively* erased once the bounce-vs-CMB-horizon-exit differential $N_{\text{tot}} - N_{\text{exit}} \gtrsim N_{\text{coh}} \sim \mathcal{O}(\text{few})$ (where N_{coh} is the contraction-phase coherence window of the matter-bounce mode functions), since the SPHEREEx accessible wavenumbers $k \sim 10^{-4} - 10^{-1} h/\text{Mpc}$ are pushed deep inside the inflationary subhorizon (bounce-era *physical* scales $k_{\text{bounce}}^{\text{phys}} \sim k_{\text{SPHEREEx}}^{\text{phys}} e^{N_{\text{tot}} - N_{\text{exit}}} \sim e^{32} k_{\text{SPHEREEx}}^{\text{phys}}$ at $N_{\text{tot}} \sim 92$, $N_{\text{exit}} \sim 60$; comoving wavenumbers k are constant by definition; the rigorous bounce-vs-SPHEREEx scale ratio is the *physical* scaling above with $k_{\text{bounce}}^{\text{phys}}/k_{\text{SPHEREEx}}^{\text{phys}} \sim e^{32}$) and the surviving bispectrum signal becomes purely vacuum-inflationary rather than matter-bounce contraction-mode. The mode-history reasoning underlying this “erased by $N_{\text{tot}} - N_{\text{exit}} \gtrsim N_{\text{coh}}$ ” statement is the standard mode-transfer ledger across four epochs: (i) in the contracting matter-bounce phase the $f_{\text{NL}} = -35/16$ signal is generated in the contraction-era cubic action at comoving wavenumbers k_{gen} that exited the contracting-phase horizon at horizon-crossing time $t_{\text{cross}}^{\text{gen}}$, (ii) at the bounce the modes are mapped onto the expanding branch with their comoving k preserved (the bounce is mode-conserving for adiabatic perturbations in a non-singular bounce with a finite minimum scale factor; only physical $k_{\text{phys}} = k/a$ scale with a), (iii) inflation then redshifts the corresponding physical scale by $a_{\text{reh}}/a_{\text{bounce}} = e^{N_{\text{tot}}}$, so a

SPHEREEx-observable comoving k today corresponds to a bounce-era physical scale $k_{\text{bounce}}^{\text{phys}} = k_{\text{obs}} e^{N_{\text{tot}} - N_{\text{exit}}}$ with $N_{\text{exit}} \sim 60$ the standard CMB horizon-exit e -fold count, and (iv) reheating preserves comoving k . The differential $e^{N_{\text{tot}} - N_{\text{exit}}} \sim e^{32}$ pushes the SPHEREEx-relevant modes to bounce-era *physical* scales $e^{32} \times k_{\text{SPHEREEx}}^{\text{phys}}$, which lie deep inside the inflationary subhorizon regime at the bounce time; the surviving observable bispectrum is dominated by the vacuum-inflationary mode functions amplified between $t_{\text{cross}}^{\text{gen}}$ and today, not by the matter-bounce contraction-mode functions that sourced $f_{\text{NL}} = -35/16$. A fully quantitative transfer function tracking the suppression coefficient across the four epochs (vacuum-inflationary amplification vs matter-bounce contraction-mode survival probability) is beyond the scope of this no-go and is reserved for a companion forecast (Paper II [2]); the present argument is the scale-history bookkeeping, not the transfer-function calculation. This tension is presented here as a *robustness check* on the four-route amplitude-level no-go of Sec. IV and the 14-barrier closure of Sec. IX, not as a co-equal closure mechanism: the no-go has already closed the four amplitude routes by which minimal ECH could source dark energy, so the structural-tension argument has nothing remaining to bind against at the route-amplitude level and reads instead as an independent consistency check from the surviving matter-bounce science case (Sec. XIII). DESI DR2 evidence for equation-of-state crossing at $3.1 - 4.2\sigma$ [9] motivates dynamical-dark-energy parameterizations (including $w_0 w_a$ and quintom-class scenarios) that can unify bounce and dark energy through mechanisms outside the ECH minimal-route catalog; these non-minimal-ECH mechanisms are not addressed by the present no-go.

E. Channel-Level Closure

The 13 mechanism-class structural constraints (14 historical catalog entries) catalog the four enumerated minimal-ECH dark-energy channels and close each under the stated assumptions (Scope and Limitations paragraph, Sec. I). The surviving science case rests on the matter-bounce f_{NL} prediction, which is compatible with the ECH framework but not derived from it (Sec. XIII).

XV. CONCLUSIONS

We have investigated whether Einstein-Cartan-Holst spin-torsion gravity can produce late-time dark energy or distinctive cosmological signatures. The answer is a channel-level closure: under the stated assumptions, the 14 mechanism-class constraints (Table IV; B8 is the observational consequence of the perturbation-transparency result B14 and is retained for historical mechanism-class completeness) constrain each of the four enumerated minimal-ECH dark-energy routes. Route R1 closes at the amplitude level via a standard torsion-

elimination derivation; Routes R2–R3 close at the amplitude level under explicitly-labeled scaling ansätze; route R4 is closed instead by a naturalness / cosmological-constant fine-tuning objection (Sec. IV F) rather than by an amplitude mismatch, because a free-coupling spectator-ALP fit can simultaneously reproduce β_{obs} and ρ_{Λ} only at $m_{\theta} \sim H_0$, re-importing the CC problem rather than solving it.

Central result: perturbation transparency.—For canonical scalar field matter, the Holst sector decouples completely from all scalar and tensor perturbation equations of motion (Sec. X). Torsion vanishes at all perturbation orders; the Holst dual contraction $e^{\mu\nu\rho\sigma} R_{\mu\nu\rho\sigma}$ vanishes identically on the Levi-Civita connection ($T = 0$) by the first (algebraic) Bianchi identity $R_{\mu[\nu\rho\sigma]} = 0$ (distinct from — and not equal to — the Pontryagin density $\propto R\tilde{R}$, which involves two curvatures and is a separate topological invariant; see Sec. X footnote for the $e\wedge e\wedge R = -NY + T\wedge T$ decomposition). This is a positive structural result: it identifies the nonperturbative parity-violating channels (ALP birefringence, primordial GWs) as parity-sensitive channels (model-dependent tests of γ_{BI} only under a derived γ_{BI} -dependent photon or tensor-parity coupling).

The 13 mechanism-class constraints (14 historical catalog entries).—Systematic analysis across 7 foundations (A–G) and 6 observational branches (H, J, L, M, N, O) established 14 mechanism-class constraints (one of which, B8, is the observational consequence of the perturbation-transparency result B14 and is retained in the catalog for historical mechanism-class completeness) on the minimal ECH dark-energy parameter space (Sec. IX). These constraints are ECH-specific: other bouncing cosmologies (notably quintom scenarios) can in principle unify bounce with late-time dark energy through mechanisms outside the ECH minimal-route catalog.

Surviving tests.—Two ECH-independent class-level predictions of the broader bounce/ALP landscape survive the channel-level closure and are testable:

1. $f_{\text{NL}} = -35/16$ (matter-bounce class): SPHEREEx tests at 1.3–2.75 σ realistic significance by ~ 2028 , discriminating from inflation ($f_{\text{NL}} \approx 0.015$) and the Cuscuton bounce ($f_{\text{NL}} \approx 0$).
2. Spectator-ALP birefringence $\beta \approx 0.27^\circ$: LiteBIRD ($\sigma(\beta) \approx 0.03^\circ$, early 2030s) detects non-zero β at $\sim 9\sigma$ (a $0.27^\circ/0.03^\circ$ overall sensitivity number). The relevant model-discrimination test, however, is the differential against the prior central value $\beta_{\text{obs}} = 0.342^\circ \pm 0.094^\circ$: LiteBIRD will distinguish the spectator-ALP-derived 0.27° from the observed 0.342° at $|0.342 - 0.27|/\sqrt{0.03^2 + 0.094^2} \approx 0.072^\circ/0.0987^\circ \approx 0.73\sigma$ (this is a heuristic combining the current Planck uncertainty $\pm 0.094^\circ$ and the projected LiteBIRD uncertainty $\pm 0.03^\circ$ in quadrature; the test is dominated by the current Planck term, so under the current Planck error alone the separation is $|0.342 - 0.27|/0.094^\circ \approx 0.77\sigma$; the 0.73σ figure is therefore not a well-defined

joint-posterior significance but a rough lower bound), NOT at the naive $|0.342 - 0.27|/0.03 = 2.4\sigma$ which would ignore the prior measurement’s $\pm 0.094^\circ$ uncertainty; in other words LiteBIRD’s $\sigma(\beta) = 0.03^\circ$ will not by itself separate the spectator-ALP value from the current WMAP+Planck birefringence central value in a model-discrimination test (a future tightening of the observational central value’s uncertainty below $\sim 0.05^\circ$ would be needed for LiteBIRD-vs-current-central tension to cross 1σ). The two numbers correspond to distinct null hypotheses (zero vs. WMAP+Planck-central); the $\sim 9\sigma$ test will not by itself separate the spectator-ALP class from generic-ALP fits to the observed signal.

Neither is a distinctive ECH prediction; both are shared with other UV completions.

Known limitations.—This work does not derive the IR effective vacuum term from first principles: $w = -1$ is assumed, not derived; the birefringence prediction lacks a derived photon-torsion coupling; α/M is a phenomenological parameter; and the MCMC uses stock CAMB (not a bespoke torsion-modified Boltzmann code). Full MCMC diagnostics, ALP parameter fitting, and NaMaster pipeline validation are in the folded MCMC/reproducibility appendices (Appendix E).

Forward.—The program continues with direct tests of the bounce framework: SPHEREEx f_{NL} forecast (Paper II [2]), multi-survey anomaly catalog including the NANOGrav 15-yr free-spectrum real-KDE GPU MCMC reanalysis used here for γ_{PTA} (Paper III [48]), and galaxy chirality catalog (Paper IV [23]).

Data and Code Availability

Minimal-ECH structural calculations, Cobaya YAML configurations, galaxy-spin pipeline code, and the implementation map are publicly available at:

<https://github.com/Hubify-Projects/bigbounce/tree/main/reproducibility>

Code and data are available at the repository above; a Zenodo-archived release will pin all artifacts to the submitted-version snapshot. The frozen MCMC chains backing the cosmology tables are committed in the bundle (under `reproducibility/cosmology/frozen/`); fresh re-verification chains must be regenerated from the supplied Cobaya configurations. The NaMaster pipeline and ALP parameter fitting are documented in companion Appendix E and are not duplicated in the present bundle.

ACKNOWLEDGMENTS

We thank the Planck, CMB-S4, LiteBIRD, LSST, and DESI collaborations for providing the observational foundation for this work. We particularly acknowledge the

foundational contributions of Nikodem Popławski, Simone Mercuri, Laurent Freidel, Djordje Minic, and Tatsu Takeuchi for their fundamental derivations connecting the Barbero-Immirzi parameter to parity-violating interactions in LQG. We acknowledge Lior Shamir for providing aggregate CW/CCW galaxy spin counts for the $A(z)$ comparison.

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, with OpenAI GPT-5/o3, xAI Grok-4, and Google Gemini 2.5 used as cross-checking and adversarial internal-review models, used for systematic barrier-cataloging, symbolic and dimensional cross-checking, perturbation-gate verification, literature triage, and manuscript preparation. All scientific claims, derivations, and numerical results were verified by the author against the committed computational artifacts—source code, frozen chains, and pipeline outputs—in the public reproducibility tree [10], 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. Computational resources were self-funded (RunPod H200 and H100 instances).

Appendix A: Complete Parameter Summary

Appendix B: Dimensional Status of the Parity-Odd Operator

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

$$\begin{aligned} [\alpha/M] &= -1, & [\varepsilon^{\mu\nu\rho\sigma} e^I_\mu e^J_\nu \mathcal{F}_{IJ\rho\sigma}] &= +2 \\ \implies [\mathcal{L}_{\text{odd}}] &= +1. \end{aligned} \quad (\text{B1})$$

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 [49–51]. In minimal ECH the *only* available scale is $\Lambda \sim M_{\text{Pl}}$: there is no intermediate threshold and no assumed cancellation. The three missing mass powers are therefore *forced* to be M_{Pl}^3 — one cannot conjure a light scale from a theory that has none — and there are exactly two dimensionally-consistent ways to supply them.

Case I — coefficient dressing (local NDA reading).—The natural coefficient of the relevant operator is $c \sim M_{\text{Pl}}^{4-1} = M_{\text{Pl}}^3$, and the only static (VEV) piece of $\varepsilon e e \mathcal{F}$ that can act as a cosmological constant is the dimension-+1 torsion/topological density T , giving $\rho \sim M_{\text{Pl}}^3 \langle T \rangle$. A coherent cosmological torsion is forbidden today by

parity and isotropy ($\langle T \rangle \rightarrow 0$, hence $\rho \rightarrow 0$), while at the bounce $\langle T \rangle \sim M_{\text{Pl}}$ returns $\rho \sim M_{\text{Pl}}^4$. Neither reproduces $(\text{meV})^4$.

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

$$\rho_\Lambda^{\text{bounce}} \sim (\alpha/M) M_{\text{Pl}}^5 \sim 10^{-2} M_{\text{Pl}}^4, \quad (\text{B2})$$

a bounce-era density that must dilute by $e^{-3N_{\text{tot}}} \sim 10^{-122}$ ($N_{\text{tot}} \approx 94$) to reach $\rho_\Lambda^{\text{obs}}$. The $+1 \rightarrow +4$ mass-dimension promotion in Case II is a *heuristic dimensional argument*: absorbing the three missing mass powers into on-shell bounce curvature factors ($R \sim M_{\text{Pl}}^2$) is a dimensional assignment consistent with the single-scale power-counting logic, but it does not constitute a full field-theoretic formalization (e.g., a controlled effective-operator construction demonstrating that the specific curvature insertion is gauge-invariant, diffeomorphism-covariant, and free of operator-mixing contributions at the bounce scale). This limitation is explicitly disclosed; the NDA order-of-magnitude conclusion that the natural density sits at $\rho_\Lambda^{\text{ECH}} \sim M_{\text{Pl}}^4$ is unchanged by this caveat, since both Case I and Case II land at the same Planck-scale ceiling via the single-scale power-counting logic, regardless of which off-shell completion is adopted.

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

$$\rho_\Lambda^{\text{ECH}} \sim M_{\text{Pl}}^4 \quad (\text{up to } \mathcal{O}(1)\text{--}\mathcal{O}(10^{-2})), \quad \text{never } (\text{meV})^4. \quad (\text{B3})$$

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

Residual assumption (explicit, not fabricated).—The bound assumes single-scale EFT: $\Lambda \sim M_{\text{Pl}}$, no intermediate threshold $\mu \ll M_{\text{Pl}}$, and no exact symmetry/topological cancellation. A non-minimal UV completion introducing a new light scale $\mu \ll M_{\text{Pl}}$ in the coefficient ($c \sim \mu^a M_{\text{Pl}}^{3-a}$), or a symmetry that zeroes the M_{Pl}^4 piece and leaves a protected small remainder, could evade the estimate — but such a scale or symmetry *is* the tuning the mechanism was meant to explain, relocating rather than solving the cosmological-constant problem. Minimal ECH contains neither, which is precisely why the no-go holds *within minimal ECH*. Full closure would require a matching calculation over parity-odd ECH-compatible completions showing that any such

TABLE VI. Complete parameter summary with priors, reference values, and physical interpretations. Rows marked † cite the internal MCMC analysis of Appendix E, refereable in-document from the committed reproducibility artifacts (Appendix H 1). The γ row uses the SU(2) full-counting scheme value $\gamma_{\text{SU}(2)} \approx 0.274$ [18, 19]; the scheme range ~ 0.037 (SU(2)–DLM spread) denotes LQG area-spectrum scheme dependence, *not* a statistical error.

Parameter	Description	Prior	Reference value	Notes
<i>Fundamental theory parameters</i>				
γ_{BI}	Barbero-Immirzi	Fixed: 0.274	0.274 (scheme range ~ 0.037)	LQG area spectrum (Eq. 2)
α/M	Parity-odd coefficient	Log-flat	$\sim 10^{-21} \text{ GeV}^{-1}$	One-loop motivated
N_{tot}	Total e-folds	[60, 120] flat	≈ 92 (fitted)	Controls Ξ
<i>Cosmological parameters (companion internal MCMC†)</i>				
H_0^\dagger	Hubble constant	From Planck prior	$67.68 \pm 1.06 \text{ km/s/Mpc}$	Recovers ΛCDM
$\Delta N_{\text{eff}}^\dagger$	Radiation proxy	[-3, 3] flat	-0.020 ± 0.169 (full-tension)	Consistent with 0
σ_8^\dagger	Clustering amplitude	Derived	0.803 ± 0.008	
Ω_m	Matter density	Derived	0.308 ± 0.005	
<i>Observational channel parameters</i>				
β	ALP birefringence	—	0.27° (midpoint)	Spectator ALP, not ECH
f_{NL}	Non-Gaussianity	—	$-35/16 = -2.1875$	Matter bounce class ^a
γ_{PTA}	PTA spectral index	—	2.567 ± 0.382 (real-KDE GPU MCMC)	Bounce $\gamma = 3.0$ at $+1.13\sigma$

^a Survives only in bounce scenarios *not* invoking the $N_{\text{tot}} \approx 92$ dark-energy dilution; under that dilution it is erased at SPHEREx-accessible scales (Sec. XIV D).

completion either (i) reproduces the M_{Pl}^4 NDA estimate, or (ii) requires an added light scale/symmetry that is itself the tuning being explained; that matching calculation is left to a companion treatment. For the minimal-ECH field content this matching is in fact immediate rather than deferred: the finite operator basis fixed by the algebraic Cartan constraint and the totally-antisymmetric minimal spin current (Sec. IV, “Scope”) contains no operator of lower off-shell dimension than the +1 operator bounded above, so single-scale NDA applied to that ceiling bounds every member of the basis; only a non-minimal light scale or protected cancellation—case (ii)—evades it, and that is the tuning being explained.

Inflationary dilution ($\mathcal{D}_{\text{inf}} \sim e^{-3N_{\text{tot}}}$) then yields $\rho_\Lambda = \Xi M_{\text{Pl}}^4$ with $\Xi = (\alpha/M) M_{\text{Pl}} \mathcal{D}_{\text{inf}}$ under Eq. (B2). The genuine cosmological-constant hierarchy is $M_{\text{Pl}}^4/\rho_\Lambda^{\text{obs}} \sim 10^{19} \text{ GeV}^4/(10^{-3} \text{ eV})^4 \sim 10^{122}$, i.e. ~ 120 orders of magnitude (the bounce-scale density entering this hierarchy is $\rho_{\text{bounce}} \sim M_{\text{Pl}}^4$, not the local pseudo-density $\rho_\Lambda^{\text{bounce}} \sim 10^{-2} M_{\text{Pl}}^4$ that Eq. (B2) labels). The dilution factor required to bridge $\sim 10^{122}$ from M_{Pl}^4 down to the observed ρ_Λ is $\mathcal{D}_{\text{inf}} \sim e^{-3N_{\text{tot}}} \sim 10^{-122}$, giving $N_{\text{tot}} \approx 122 \ln 10/3 \approx 94$ e-folds (consistent at the $\sim 2\%$ level with the structural-tension $N_{\text{tot}} \approx 92$ quoted in Sec. XIV D; the small offset reflects that the structural tension uses Eq. (B2) as the input ansatz, while the genuine M_{Pl}^4 -to- $\rho_\Lambda^{\text{obs}}$ hierarchy uses the unrescaled Planck density).

Sharper dependency statement.—The precise value $N_{\text{tot}} \approx 92$ vs the independent M_{Pl}^4 -to- $\rho_\Lambda^{\text{obs}}$ estimate $N_{\text{tot}} \approx 94$ (a $\sim 2\%$ offset) *does* depend on the ansatz choice in Eq. (B2); however, the *overall scale separation* between $\rho_\Lambda^{\text{bounce}}$ and $\rho_\Lambda^{\text{obs}}$ is ~ 120 orders of magnitude under *any* non-cancellation assumption (whether the bounce-scale density is M_{Pl}^4 , $10^{-2} M_{\text{Pl}}^4$, or $10^{-4} M_{\text{Pl}}^4$ shifts the required e-fold count by $\mathcal{O}(\text{a few})$, not by orders

of magnitude), and this scale separation is what drives the 13 mechanism-class structural barriers (Sec. IX). The barriers close the bounce-to-dark-energy route on amplitude-budget and operator-counting grounds whose qualitative conclusions are ansatz-independent; only the $\sim 2\%$ precision of the headline N_{tot} figure is ansatz-dependent. Readers should therefore treat $N_{\text{tot}} \approx 92$ as an order-of-magnitude estimate ($N_{\text{tot}} = 92 \pm 2$ accounting for the ansatz-choice systematic), not as a precise number; the broader structural closure does not depend on the precise e-fold value.

We emphasize that the present treatment labels the on-shell scaling explicitly as a phenomenological ansatz rather than as a derivation: a controlled EFT-level construction (an operator-basis closure for a local dimension-+4 parity-odd interaction with all required M_{Pl} factors in the coupling coefficient) remains left to a separate companion treatment.

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

The torsion-elimination step (Sec. II A 2, Eqs. (3)–(14)) generates, at the single prefactor $\kappa = 8\pi G = M_{\text{Pl}}^{-2}$, the axial-axial contact term $c_{AA}(J^5 \cdot J^5)$ and—at finite Barbero-Immirzi parameter—the vector-axial Holst partner $c_{VA}(J \cdot J^5)$. We prove that the complete Fierz rearrangement of these operators closes onto the enumerated dimension- ≤ 6 parity-relevant basis with *no* escape operator and *no* change of M_{Pl} power.

On the 16-dimensional Clifford basis, grouped into the five Lorentz classes $(S, V, T, A, P) = (\mathbf{1}, \gamma^\mu, \sigma^{\mu\nu}, \gamma^\mu \gamma^5, \gamma^5)$, Fierz rearrangement of a current-current product is the linear map $F_{AB} = \frac{1}{4} \text{Tr}_{\text{class}}(\Gamma^A \Gamma_B \Gamma_A \Gamma^B)/d_B$, $d_B = \text{Tr}(\Gamma^B \Gamma_B)$.

Computed from the explicit Dirac matrices in the paper's mostly-plus signature (script `arxiv/scripts/fierz_lemma_check.py`), it equals the standard matrix [52, 53]

$$F = \frac{1}{4} \begin{pmatrix} 1 & 1 & 1 & 1 & 1 \\ 4 & -2 & 0 & 2 & -4 \\ 6 & 0 & -2 & 0 & 6 \\ 4 & 2 & 0 & -2 & -4 \\ 1 & -1 & 1 & -1 & 1 \end{pmatrix}, \quad F^2 = \mathbb{1}, \quad (\text{C1})$$

in the order (S, V, T, A, P) ; the involution $F^2 = \mathbb{1}$ is verified symbolically. Applying F to the generated operators gives the exact decompositions

$$(J^5 \cdot J^5) \xrightarrow{\text{Fierz}} \frac{1}{4} SS + \frac{1}{2} VV - \frac{1}{2} AA - \frac{1}{4} PP, \quad (\text{C2})$$

$$(J \cdot J) \xrightarrow{\text{Fierz}} \frac{1}{4} SS - \frac{1}{2} VV + \frac{1}{2} AA - \frac{1}{4} PP, \quad (\text{C3})$$

while the parity-odd $(J \cdot J^5)$ Holst partner rotates only within the $\{V, A\}$ block ($F_{VA} = F_{AV} = \frac{1}{2}$), remaining a dimension-6 parity-odd cross term. 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 $\kappa = M_{\text{Pl}}^{-2}$ prefactor exactly. This is the explicit Fierz-by-Fierz projection lemma: within single-species minimal ECH the generated four-fermion tower is Fierz-closed and uniformly M_{Pl}^{-2} -suppressed, so the single-scale NDA ceiling of App. B that bounds one representative bounds the entire finite basis. (For multiple fermion species the Lorentz-class closure and the M_{Pl} -power argument are unchanged—Fierz acts on Dirac structure, not flavor—and only the flavor-off-diagonal coefficient bookkeeping is added.) The only operators evading this bound require a *non-minimal* completion (a new light scale $\mu \ll M_{\text{Pl}}$, trace/tensor torsion irreps, or a dynamical Chern–Simons field), which is the stated scope boundary of the no-go.

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

Appendix D: Line-of-Sight Birefringence from the Maxwell–Chern–Simons Operator

This appendix derives the rotation-angle mapping $\beta = (\alpha/2M) \Delta\phi = (\alpha f_a/2M) \Delta\theta$ used in Eq. (19), fixing the factor of 1/2 from first principles in the operator normalization used throughout this paper.

Setup.—For the spectator-ALP benchmark used in Sec. IV F, assume the Maxwell–Chern–Simons form

$$\mathcal{L} \supset -\frac{1}{4} F_{\mu\nu} F^{\mu\nu} - \frac{1}{4} \frac{\alpha}{M} \phi F_{\mu\nu} \tilde{F}^{\mu\nu}, \quad (\text{D1})$$

with $\tilde{F}^{\mu\nu} = \frac{1}{2} \epsilon^{\mu\nu\rho\sigma} F_{\rho\sigma}$ and ϕ the dim+1 canonical pseudoscalar (homogeneous: spatial gradients negligible for the cosmological background field considered here) [54, 55]. Throughout this appendix we use the dimensionless angle $\theta \equiv \phi/f_a$ for evolution equations and quote $\Delta\theta$ or $\Delta\phi = f_a \Delta\theta$ interchangeably for the line-of-sight excursion; the single-convention bridge $\phi = f_a \theta$ is recorded in the footnote at the operator's first appearance in Sec. IV F.

Field equation in FRW.—Varying θ 's canonical action with potential $V = m_\theta^2 f^2 (1 - \cos\theta)$ in a flat FRW background gives

$$\ddot{\theta} + 3H\dot{\theta} + m_\theta^2 \sin\theta = 0, \quad (\text{D2})$$

the equation integrated numerically by the released pipeline (`research/branch_R_alp_birefringence/phase2_mcmc/alp_ode.py`; frozen initial condition $\dot{\theta}(t_{\text{init}}) = 0$ at $z_{\text{init}} = 3000$).

Helicity dispersion.—In conformal time ($ds^2 = a^2(-d\eta^2 + d\vec{x}^2)$) the Maxwell sector of Eq. (D1) is conformally invariant, and the $\theta F\tilde{F}$ term, being topological up to the θ prefactor, carries no additional factors of a . In Coulomb–temporal gauge ($A_0 = 0$, $\partial_i A^i = 0$; because $\partial_\mu \theta = \theta' \delta_\mu^0$ for the homogeneous background field, the parity-odd term contributes only through the spatial $\epsilon^{ijk} A_j \partial_k$ structure and the longitudinal mode stays non-dynamical), decomposing the transverse field into circular-polarization modes $A_\pm(\eta, k)$, the modified wave equation reads

$$A''_\pm + \left[k^2 \mp \frac{\alpha}{M} \phi' k \right] A_\pm = 0, \quad (\text{D3})$$

where $' \equiv d/d\eta$, k is comoving, and $\phi' \equiv d\phi/d\eta = f_a \theta'$; with $[\alpha/M] = -1$, $[\phi'] = +2$, each bracketed term carries dimension +2 as required for consistency with k^2 . For $k \gg (\alpha/M)\phi'$ (satisfied by ~ 30 orders of magnitude for CMB photons against the cosmological field considered here: $(\alpha/M)\phi' \sim 10^{-30} \text{ eV}^{-1} \times [\mathcal{O}(1) f_a] \times H_0 \sim 10^{-35} \text{ eV}$ for an $\mathcal{O}(1)$ dimensionless excursion $\Delta\theta$ accumulated over a Hubble time at $f_a \sim M_{\text{Pl}}$ and the cosmological-tuning condition $m_\theta \sim H_0$ that sets $\dot{\theta} \sim H_0$ (in conformal time, with $\phi' \equiv d\phi/d\eta$), versus a CMB photon $k \sim 6 \times 10^{-4} \text{ eV}$ at 150 GHz, a ratio of $\sim 10^{31}$) the WKB phase of each helicity is $\int \omega_\pm d\eta$ with $\omega_\pm \simeq k \mp \frac{1}{2}(\alpha/M)\phi'$.

Rotation angle.—A linear polarization is a fixed-phase superposition of the two helicities; its position angle rotates by half the accumulated helicity phase difference:

$$\begin{aligned}\beta &= \frac{1}{2} \int_{\eta_{\text{em}}}^{\eta_{\text{obs}}} (\omega_- - \omega_+) d\eta = \frac{1}{2} \frac{\alpha}{M} \int_{\eta_{\text{em}}}^{\eta_{\text{obs}}} \phi' d\eta \\ &= \frac{\alpha}{2M} [\phi(\eta_{\text{obs}}) - \phi(\eta_{\text{em}})] \equiv \frac{\alpha}{2M} \Delta\phi = \frac{\alpha f_a}{2M} \Delta\theta.\end{aligned}\tag{D4}$$

The result is achromatic (no k dependence), independent of the expansion history except through the endpoint values of ϕ (equivalently the dimensionless angle $\theta = \phi/f_a$; a total derivative along the line of sight — the well-known “ $2\beta = g \Delta\phi$ ” property [55]), and exact at leading WKB order. This is Eq. (19)’s mapping; with the spectator-ALP identification $\alpha/M \equiv C_{a\gamma} \alpha_{\text{em}}/(2\pi f_a)$ (cf. the basis-conversion footnote at the first quote of $\alpha/M = 10^{-21} \text{ GeV}^{-1}$ in Sec. IV F for the $1/(2\pi)$ -vs- $1/(4\pi)$ normalization gap) and $\Delta\theta = \Delta\phi/f_a$ it reproduces the companion pipeline’s $\beta = (\alpha_{\text{em}} C_{a\gamma}/4\pi)(\Delta\phi/f_a)$ (the convention block of the companion’s §VI⁸), closing the normalization chain from the Lagrangian to the numerical prediction; this mapping is what Eq. (19) and the Route-2 estimates of Sec. IV use. The overall sign of β matches the WMAP+Planck convention of Eskilt–Komatsu, in which a freely rolling field with $\theta(\eta_{\text{obs}}) > \theta(\eta_{\text{em}})$ produces $\beta > 0$.

Appendix E: Cosmological Constraint Methodology: $\Lambda\text{CDM}+\Delta N_{\text{eff}}$ MCMC Proxy and the ECH-Sector ΔN_{eff} Bound

This appendix documents the stock-CAMB $\Lambda\text{CDM}+\Delta N_{\text{eff}}$ MCMC proxy run, the first-principles bespoke ECH-sector ΔN_{eff} bound, and the full cosmological fits and model comparison. These analyses are numerical consistency cross-checks that contextualize the main-text no-go; none tests or verifies the spin-torsion sector. All quantitative values here are reproducible from the committed artifacts (Appendix H 1).

1. Cosmological Tensions: H_0 and σ_8

The bounce scenario motivates extending ΛCDM by ΔN_{eff} (particle production at the bounce) as a phenomenological proxy parameter; $(\omega/H)_0$ (angular momentum transfer) and Ω_k are fixed to zero in the actual sampled MCMC configuration of this paper (per

the explicit parameter-scope clarification in Sec. E 3 a; the $(\omega/H)_0$ parameter is discussed in the main text as a phenomenological bounce-class indicator but is not separately sampled here). The full-tension dataset combination includes the SH0ES [56] M_B -anchor likelihood (H0.riess2020Mb, entering through the supernova absolute magnitude rather than as a direct H_0 prior) in the Cobaya likelihood configuration, but the high-precision Planck NPIPE (PR4) CamSpec high- ℓ TT-TEEE + Planck 2018 low- ℓ TT/EE + Planck 2018 lensing likelihoods carry sufficient inverse-variance weight that the posterior H_0 in the proxy run is pulled to 67.68 ± 1.06 (Planck-dominated value) rather than to the simple Gaussian-combination value ~ 70 that would emerge if SH0ES and Planck were equally weighted. We do not therefore claim that the SH0ES tension is resolved or even moved by adding ΔN_{eff} in stock CAMB; in this configuration the ΔN_{eff} posterior is consistent with zero and the H_0 posterior is dominated by Planck. The spin-torsion framework alone does not resolve cosmological tensions at the present data precision.

2. Stock-CAMB $\Lambda\text{CDM}+\Delta N_{\text{eff}}$ MCMC: Generic Radiation-Proxy Test (Not a Spin-Torsion Theory Module)

Scope statement.—“MCMC verification” refers throughout to a stock CAMB run of ΛCDM extended by ΔN_{eff} as a free parameter. *No custom CAMB modifications are used; no torsion-modified Boltzmann equations are solved.* The MCMC therefore tests whether the data prefer an extra radiation-like degree of freedom, treated as a generic phenomenological proxy for the spin-torsion sector’s possible effective radiation contribution. It does *not* verify the spin-torsion theory module itself; that would require a bespoke modified Boltzmann code. What the MCMC *does* anchor, via the bespoke ECH-sector calculation in Sec. E 2 a below, is a conservative observational envelope on the actual radiation contribution of the minimal Einstein–Cartan–Holst spin-torsion sector, which that calculation shows is Planck-suppressed and therefore far smaller than the proxy prior permits.

a. Bespoke ECH-sector ΔN_{eff} from the torsion-induced four-fermion interaction

Rather than leave ΔN_{eff} as a purely generic radiation parameter, we compute the *actual* contribution of the minimal Einstein–Cartan–Holst (ECH) spin-torsion sector to the early-universe effective number of relativistic species. In minimal ECH gravity the torsion tensor is *algebraic* (non-propagating; it carries no kinetic term) and is sourced purely by the fermion axial current $A^\mu = \bar{\psi}\gamma^5\gamma^\mu\psi$ through the Cartan constraint [12]. Integrating torsion out of the action leaves an effective

⁸ The WKB condition $k \gg (\alpha/M)\phi'$ entering this appendix is satisfied by ~ 30 orders of magnitude: $(\alpha/M)\phi' \sim 10^{-35} \text{ eV}$ vs. $k \sim 6 \times 10^{-4} \text{ eV}$ for a 150 GHz CMB photon (see computation block above Eq. (D3)). This recomputation is self-contained within App. D; no companion result is required to establish the WKB approximation.

four-fermion contact interaction of Nambu–Jona-Lasinio type [16, 57],

$$\mathcal{L}_{4f} = -\frac{3\kappa^2}{16} \frac{\gamma^2}{\gamma^2 + 1} (\bar{\psi}\gamma^5\gamma^\mu\psi)^2, \quad \kappa^2 = 8\pi G_N = M_{\text{Pl}}^{-2}, \quad (\text{E1})$$

where γ is the Barbero–Immirzi parameter of the Holst term [16, 42]; in the pure Einstein–Cartan–Sciama–Kibble limit $\gamma \rightarrow \infty$ the bracket $\rightarrow 1$ and Eq. (E1) reduces to the canonical $-\frac{3}{2}\pi G_N (\bar{\psi}\gamma^5\gamma^\mu\psi)^2$ [57]. The coefficient is fixed by the geometry, not fitted, so this is genuine ECH-sector physics rather than a generic parametrization.

The crucial feature of Eq. (E1) is that the coupling is $\propto \kappa^2 = 8\pi G_N \sim M_{\text{Pl}}^{-2}$: the interaction is Planck-suppressed. In a thermal fermion gas the coherent spin-torsion energy density is $\rho_{\text{tor}} \sim \alpha n_f^2$ with $\alpha = \kappa^2(\hbar c)^2/32$ and number density $n_f \propto T^3$ for relativistic species [58]. Explicitly, the energy density contributed by the dimension-6 contact operator of Eq. (E1) is the thermal expectation value $\rho_{\text{tor}} \sim \kappa^2 \langle (\bar{\psi}\gamma^5\gamma^\mu\psi)^2 \rangle_T$; for a relativistic plasma in kinetic equilibrium the coherent axial-current bilinear factorizes at leading order into the square of the (spin-summed) fermion number density, $\langle (\bar{\psi}\gamma^5\gamma^\mu\psi)^2 \rangle_T \sim n_f^2 \sim (g_* T^3)^2$ (a parametric estimate: the standard finite-temperature scaling of a four-fermion (Fermi-type) contact operator in the early-universe plasma [59], with $\mathcal{O}(1)$ coefficients dropped rather than a rigorous NJL finite- T computation; *a full finite-temperature NJL computation of the exact prefactor is outside the scope of this companion, but is observationally moot: the 10^{-44} BBN suppression of $\Delta N_{\text{eff}}^{(\text{ECH})}$ renders any $\mathcal{O}(1)$ correction factor physically irrelevant at the \gtrsim forty-orders-of-magnitude margin from CMB- S_4 sensitivity), so that*

$$\rho_{\text{tor}} \sim G_N T^6, \quad \frac{\rho_{\text{tor}}}{\rho_{\text{rad}}} \sim G_N T^2 = \left(\frac{T}{M_{\text{Pl}}}\right)^2, \quad (\text{E2})$$

since $\rho_{\text{rad}} \sim T^4$. The spin-torsion energy density thus redshifts as a^{-6} (stiff-fluid-like), faster than radiation (a^{-4}), and becomes dynamically relevant only at densities approaching the Planck/bounce scale [58]. Its contribution to the effective radiation budget at the epochs probed by the data is

$$\Delta N_{\text{eff}}^{(\text{ECH})} \sim \left(\frac{T}{M_{\text{Pl}}}\right)^2 \sim \begin{cases} 1.7 \times 10^{-43}, & T \sim 1 \text{ MeV}, \\ 1.1 \times 10^{-56}, & T \sim 0.26 \text{ eV}, \end{cases} \quad (\text{E3})$$

at BBN and recombination respectively, evaluated with the reduced Planck mass $M_{\text{Pl}} = (8\pi G_N)^{-1/2} \simeq 2.44 \times 10^{18} \text{ GeV}$ used throughout (and in Eq. E2), so the boxed numbers are reproducible directly from $(T/M_{\text{Pl}})^2$. This is the bespoke ECH-sector result: *the minimal spin-torsion sector predicts a strictly negligible ΔN_{eff} in the radiation era*, some forty orders of magnitude below any foreseeable CMB or BBN sensitivity [e.g. CMB- S_4 targets $\sigma(N_{\text{eff}}) \sim 0.03$ [60]]. No published analysis finds a

measurable torsion contribution to N_{eff} in the radiation era; the nearest dedicated BBN study constrains a distinct torsion-induced neutrino spin-flip channel, not the four-fermion energy density, and likewise reports no detectable effect [61]. We therefore do *not* claim a specific non-zero torsion ΔN_{eff} ; the honest, leading-parametric-order first-principles estimate yields the upper bound of Eq. (E3).

This reframes the MCMC that follows. The stock-CAMB $\Lambda\text{CDM} + \Delta N_{\text{eff}}$ run is *not* a proxy standing in for an unknown ECH prediction; it is a *conservative observational envelope* on the bespoke ECH contribution derived above. The data-driven 95% bound $\Delta N_{\text{eff}} \lesssim 0.31\text{--}0.40$ (Sec. E3) exceeds the theoretical ECH value Eq. (E3) by more than forty orders of magnitude, so the ECH spin-torsion sector is trivially consistent with the data through this channel, and the null-consistency of the run is exactly what the ECH physics predicts. The one place the torsion four-fermion term is *not* negligible—within a Hubble time of the bounce, where $T \rightarrow M_{\text{Pl}}$ and the ratio in Eq. (E2) is $\mathcal{O}(1)$ —is a genuinely non-perturbative regime outside the linear-perturbation reach of any Boltzmann code (stock or bespoke) and outside the scope of this companion.

The proxy run (Cobaya v3.6.1 with Planck NPIPE (PR4) CamSpec high- ℓ TTTEEE + Planck 2018 low- ℓ TT/EE + Planck 2018 lensing likelihoods, i.e. `planck_NPIPE_highl_CamSpec.TTTEEE + planck_2018_lowl.TT + planck_2018_lowl.EE + planck_2018_lensing.clik` in the Cobaya YAML) has produced two frozen dataset combinations with publication-quality convergence (limitation: the PR4/NPIPE high- ℓ + 2018-release low- ℓ /lensing mixture is the standard Cobaya pairing; the c15 verification re-run uses `planck_2020_lollipop.lowlE` and `planckpr4lensing` in place of the frozen-chain `planck_2018_lowl.EE` and `planck_2018_lensing.clik`, making c15 a release-pairing robustness rerun; the 0.04σ agreement in ΔN_{eff} across this likelihood substitution provides an empirical bound on pairing-induced bias at the quoted precision; see Sec. E3 *Release-pairing note*), plus a third Planck-only run currently at sub-convergence sample count and not aggregated into any frozen-posterior summary statistic in this paper. Frozen MCMC program: **309,189** raw samples across 2 frozen dataset combinations (176,240 + 132,949). We report this as a $\Lambda\text{CDM} + \Delta N_{\text{eff}}$ null-consistency test: the data are consistent with $\Delta N_{\text{eff}} = 0$ in stock CAMB. Sampling configuration (per the on-disk Cobaya YAMLs): ΔN_{eff} enters as `nnu` with a flat prior $N_{\text{eff}} \in [2.046, 5.046]$ (i.e. $\Delta N_{\text{eff}} \in [-1, +2]$); the neutrino mass is held at the CAMB default $\Sigma m_\nu = 0.06 \text{ eV}$ with one massive eigenstate (`num_massive_neutrinos=1`), and Y_{He} follows the CAMB BBN-consistent default (no explicit override in any of the four YAMLs). Specifically, CAMB uses its PArthENoPE-derived BBN-consistency module by default (explicit Cobaya/CAMB flag: `bbn_predictor: 'PArthENoPE'`, declared at the

theory.camb.extra_args block of each `cobaya_*.yaml`, pinning the BBN predictor for absolute reproducibility): Y_{He} is set self-consistently with N_{eff} at each MCMC sample. The prior range $N_{\text{eff}} \in [2.046, 5.046]$ remains within the calibrated domain of the CAMB BBN module; no free- Y_{He} control run was performed, as the default BBN-consistent track is the standard choice for this type of proxy analysis. The prior allows negative ΔN_{eff} (down to -1), so the results are two-sided posterior means, not physical “extra-species” limits. A physical extra-radiation degree of freedom cannot contribute $\Delta N_{\text{eff}} < 0$; imposing that boundary therefore changes the reported statistic from a two-sided mean (which can sit at a mildly negative value) to a one-sided upper bound read off the posterior renormalised on the physical half-line $\Delta N_{\text{eff}} \geq 0$. Under that physically motivated restriction $\Delta N_{\text{eff}} \geq 0$, the one-sided 95% upper limit from the full-tension chain is $\Delta N_{\text{eff}} < 0.31$, defined precisely as the 95th percentile of the posterior *renormalised on $\Delta N_{\text{eff}} \geq 0$* (we discard the $\Delta N_{\text{eff}} < 0$ tail of weighted samples and rescale the surviving weights so $\int_0^\infty p(\Delta N_{\text{eff}}) d\Delta N_{\text{eff}} = 1$, then read the 95th percentile of the renormalised CDF; this post-processing strategy is conservative when the unconstrained mode is mildly negative, as is the case here); the `planck+BAO+SN` chain gives $\Delta N_{\text{eff}} < 0.40$ under the same definition (direct weighted-sample computation on the committed chain with 30% burn-in: 93,066 post-burnin samples, of which $\sim 62\%$ satisfy $\Delta N_{\text{eff}} \geq 0$, and the 95th percentile of the renormalised truncated CDF reads 0.4012). These bounds are consistent with the two-sided means quoted throughout; both conventions are reported because a published-extra-species interpretation of the run requires the one-sided limit.⁹

⁹ Sample-count stratification (reconciliation): **309,189** is the sum of the two frozen combinations (176,240 + 132,949 raw accepted samples). After removing the first 30% of each chain as burn-in, $176,240 \times 0.7 + 132,949 \times 0.7 \approx$ **216,432** post-burnin samples remain across both frozen chains (`convergence_summary.json`). For the full-tension subset specifically, $176,240 \times 0.7 \approx$ **123,368** post-burnin (the **119,617** figure in Fig. 8 reflects additional `getdist` effective-sample weight-based thinning of this subset only). The post-burnin count of the full-tension subset alone is 123,129 (within $\pm 1\%$ of the 123,368 exact computation, with the small offset reflecting the chain-end-truncation of partial samples at the burn-in cut); the correct both-chains post-burnin total is 216,432. *Burn-in reconciliation note:* The frozen `planck_bao_sn_20260312.1954_convergence_report.txt` reports “Burn-in: 20%” and post-burnin samples = 106,361, reflecting the per-chain GetDist readout that averages the parallel chains (`burn_in: 0.1` in their `cobaya_config.yaml`) with the original chain (`burn_in: 0.3`). The 30% figure used throughout this paper is conservative, matches the original-chain configuration documented in `COUNT_EXPLANATION.md`, and gives 93,064 post-burnin for this subset (the frozen `parameter_summary_CORRECTED.json` records 93,066 for the same cut; the 2-sample difference reflects per-chain rounding of the 30% burn-in index at the chain ends and is immaterial to all quoted statistics). The 106,361 figure at 20% is the value reported by GetDist (`getdist.MCSamples` loaded

a. Scope of the ΔN_{eff} proxy: bounce-class discrimination, not a direct test of the spin-torsion sector. We emphasize that the stock-CAMB $\Lambda\text{CDM}+\Delta N_{\text{eff}}$ proxy run does *not* test the ECH spin-torsion sector directly. The Hehl-Datta-Mercuri parity-even four-fermion contact interaction that survives torsion elimination [12, 42] is dimension-6 and M_{Pl}^{-2} -suppressed¹⁰: its leading Boltzmann effect is a scattering-amplitude shift, not a relativistic species, and it does *not* produce a ΔN_{eff} at recombination. The *minimal* matter-bounce class [1] (here “minimal” = the single-scalar $w = 0$ matter-dominated contraction phase, distinguished from the broader matter-bounce family including ekpyrotic, cyclic, Cuscuton, string-gas, and quintom-bounce variants) predicts $\Delta N_{\text{eff}} \approx 0$ by construction (no light bounce-internal species are thermalized at recombination); the proxy run *confirming* $\Delta N_{\text{eff}} = -0.020 \pm 0.169$ (full-tension) and $+0.058 \pm 0.179$ (Planck+BAO+SN) is therefore consistent with the minimal matter-bounce prediction, but the absence of a ΔN_{eff} detection is not a discriminator between minimal-ECH and standard ΛCDM at the present data precision. We frame the proxy as a bounce-class *compatibility check* (minimal matter-bounce-class scenarios without prolonged post-bounce inflation predict $\Delta N_{\text{eff}} \approx 0$), not as a posterior-preference test against a competing model.

M_B - H_0 joint-posterior offset check. The joint posterior mean ($M_B = -19.263$, $H_0 = 67.68$) is fully consistent with an active `sn.pantheonplus` likelihood: `sn.pantheonplus` enforces a soft constraint on the dimensionally consistent combination $M_B - 5 \log_{10}(h) \approx \text{const}$, with $h = H_0/(100 \text{ km s}^{-1} \text{ Mpc}^{-1})$ (equivalently $M_B - 5 \log_{10}(H_0/[\text{km s}^{-1} \text{ Mpc}^{-1}]) + 10 = \text{const}$), along the SN distance-modulus degeneracy. At the SHOES [56] anchor ($M_B = -19.253$, $h = 0.7304$), the constant is $-19.253 - 5 \log_{10}(0.7304) = -18.571$. At the chain joint mean ($M_B = -19.263$, $h = 0.6768$), the constant is $-19.263 - 5 \log_{10}(0.6768) = -18.415$, i.e. a 0.156 mag offset from the Riess anchor along the Pantheon+ constraint axis. This offset is $\sim 3.2\sigma$ relative to the chain’s

from the `planck_bao_sn_20260312.1954` freeze with `ignore_rows = 0.20` averaging the per-chain configurations documented above); the paper’s 216,432 combined total uses the conservative 30% cut uniformly for both frozen chains. The third (Planck-only) dataset combination (114,992 raw samples; $\hat{R} - 1 \sim 0.05$) is still accumulating samples, is *not* reported in Table VII (which contains only the two frozen combinations), and is *not* aggregated into the 309,189-sample headline anywhere in this paper.

¹⁰ The dimension-6 four-fermion contact operator is interpreted here in the low-energy effective field theory (EFT) below the strong-coupling / torsion-resolution scale $\Lambda_{\text{strong}} \sim M_{\text{Pl}}/\sqrt{\gamma_{\text{BI}}}$ set by the inverse Barbero-Immirzi parameter γ_{BI} (cf. [42], where the contact interaction is derived; the strong-coupling scale itself is a heuristic EFT-validity estimate); above this scale the contact-operator description breaks down and the full Holst-sector dynamics with propagating torsion must be retained. All recombination-era and late-time observables addressed in this paper sit at energies $E \ll \Lambda_{\text{strong}}$, where the EFT treatment is controlled.

TABLE VII. Λ CDM+ ΔN_{eff} proxy MCMC results from frozen Cobaya v3.6.1 chains (CAMB v1.6.5, stock; an earlier pilot at Cobaya v3.5 reproduced the same posterior means to within 0.1σ , so software-version-induced shifts on ΔN_{eff} or H_0 at sub- 0.1σ precision are not formally characterized here but are not load-bearing for any quoted result; *no torsion modifications*). All values are posterior means $\pm 1\sigma$. Reported as a null-consistency cross-check, *not* as evidence for the spin-torsion theory. Under the physically motivated restriction $\Delta N_{\text{eff}} \geq 0$ (posterior renormalised on the non-negative half and 95th percentile read from the truncated CDF), the one-sided 95% upper limits are: full-tension $\Delta N_{\text{eff}} < 0.31$; Planck+BAO+SN $\Delta N_{\text{eff}} < 0.40$ (see Sec. E 2 for the exact definition and chain-recomputed derivation). **Extraction note:** the numerical values in this table are read from `parameter_summary_CORRECTED.json`; the legacy `parameter_summary.json` export carried an off-by-one column-index bug and has been deleted from the public repository (see the Column-permutation warning in Appendix H 1). $S_8 \equiv \sigma_8 (\Omega_m/0.3)^{1/2}$, computed as a Cobaya derived parameter with exactly this definition. The full-tension column includes the DES-Y3 S_8 Gaussian prior 0.776 ± 0.017 (`cobaya_full_tension.yaml`). The Planck+BAO+SN S_8 marginal is 0.827 ± 0.010 from a direct GetDist pass over the frozen chains (132,949 samples). With the chain-recomputed marginal, the full-tension posterior $S_8 = 0.814 \pm 0.008$ is consistent with the naive inverse-variance-weighted combination (denoted \otimes) of the Planck+BAO+SN marginal and the DES-Y3 prior ($0.827 \pm 0.010 \otimes 0.776 \pm 0.017 = 0.814 \pm 0.009$; agreement at the 0.01σ level). The Planck+BAO+SN marginal sits in 2.6σ two-Gaussian tension with DES-Y3 (posterior-overlap integral $\int \min(p_1, p_2) dS_8 = 0.05$, where p_1 and p_2 are the two one-dimensional S_8 marginal densities evaluated on a common $\Delta S_8 = 10^{-4}$ trapezoidal grid spanning $S_8 \in [0.70, 0.90]$); the full-tension posterior sits at 2.0σ from the DES-Y3 prior (overlap 0.12 on the same grid); this 2.0σ residual is a within-stack readout because the DES-Y3 Gaussian prior is itself an active likelihood in the full-tension chain, so the full-tension S_8 is partially anchored to DES-Y3 and the number is not a measurement-vs-measurement tension but a within-stack posterior shift; overlay artifact `reproducibility/cosmology/c13_s8_desy3_overlay.json`.

Parameter	Full-tension	Planck+BAO+SN
H_0 [km s $^{-1}$ Mpc $^{-1}$]	67.68 \pm 1.06	67.78 \pm 1.09
ΔN_{eff}	-0.020 \pm 0.169	+0.058 \pm 0.179
σ_8	0.803 \pm 0.008	0.812 \pm 0.009
S_8	0.814 \pm 0.008	0.827 \pm 0.010
Ω_m	0.308 \pm 0.005	0.312 \pm 0.006
τ	0.054 \pm 0.007	0.056 \pm 0.007
n_s	0.965 \pm 0.006	0.967 \pm 0.006
Chains	6	6
Total samples	176,240	132,949
Worst $\hat{R} - 1^a$	0.001	0.003
Min ESS	4,744	4,692

^a Worst row is n_s in the full-tension combination at $\hat{R} - 1 = 9.74 \times 10^{-4}$; all sampled parameters across both frozen combinations satisfy $\hat{R} - 1 \leq 3 \times 10^{-3}$. The full-tension chain samples 17 parameters: 7 cosmological + 9 Planck likelihood nuisance (A_{planck} , amp_{143} , amp_{217} , $\text{amp}_{143 \times 217}$, n_{143} , n_{217} , $n_{143 \times 217}$, calTE , calEE) + 1 supernova absolute magnitude (M_b , entering via the Pantheon+ SN and SH0ES `H0.riess2020Mb` likelihoods rather than the Planck likelihood stack). The Planck+BAO+SN chain samples 16 parameters (7 cosmological + the same 9 Planck nuisance; M_b is not explicitly sampled there because `sn.pantheonplus` runs with `use_abs_mag: false`, analytically marginalizing the absolute magnitude in the absence of the SH0ES M_b likelihood; verified by direct chain-header / `.updated.yaml` blocking-list inspection). Full per-parameter convergence tables (\hat{R} , ESS, drift, per dataset) are archived at `reproducibility/cosmology/convergence_latest.csv` and mirrored in the HuggingFace chain-diagnostics dataset (Appendix H 1).

$\sigma_{M_B} = 0.049$ marginal width and is the same Hubble tension manifesting in the M_B axis (3.2σ in chain- σ units, versus the canonical 3.6σ when the tension is expressed in distance-ladder terms in the H_0 axis; the two figures are derived from different estimators — chain-marginal width in M_B vs survey-vs-survey H_0 comparison — and are therefore *not directly comparable* as single-number tensions). We note this 3.2σ figure is a descriptive offset measure normalized by the marginal σ_{M_B} only: it does not condition on the M_B - H_0 covariance along the SN degeneracy direction, nor on the uncertainty of the Pantheon+ constraint itself, and is therefore not a properly conditioned tension statistic — the conditioned statement remains the canonical H_0 -axis tension. The chain posterior is therefore the correct compromise of three inputs: the Riess M_B -prior preferring $M_B = -19.253$ ($\sigma = 0.027$); the Pantheon+ likelihood preferring $M_B - 5 \log_{10} h = -18.571$ along the SN degeneracy;

and the Planck-CMB-likelihood preferring $H_0 \approx 67.5$ via the acoustic-scale calibration. With ΔN_{eff} bounded near zero by the joint data, the model cannot resolve the tension and the chain settles into the 3.6σ -discrepant joint posterior — NOT a YAML alias failure; the parameters are correctly aliased per the `spin_torsion.input.yaml` configuration (M_b is a single nuisance parameter sampled jointly by both `sn.pantheonplus` and `H0.riess2020Mb`).

Key finding.—Both frozen datasets find ΔN_{eff} consistent with zero and H_0 consistent with Planck Λ CDM at 0.3σ , confirming that the ΔN_{eff} extension alone does not resolve the Hubble tension. Current data neither require nor exclude a small positive phenomenological ΔN_{eff} proxy for the bounce-class effective radiation contribution; CMB-S4 ($\sigma(N_{\text{eff}}) \sim 0.03$ [60]) will sharpen this constraint substantially relative to the current Planck+BAO+SN level.

Independent cross-validation.—Liu *et al.* [43] con-

strained an EC torsion model using DESI DR2 [9] + Pantheon+ [62] + DES-SN5YR [63] + Planck 2018, finding torsion preferred by AIC ($\Delta\text{AIC} = -5.7$ to -6.6) but with the torsion parameter itself consistent with zero ($\alpha = -0.00066 \pm 0.00098$). Their headline values $H_0 = 68.41 \pm 0.32$ km/s/Mpc and $S_8 = 0.812 \pm 0.006$ agree with our Planck+BAO+SN chain at 0.5σ in H_0 ($|67.78 - 68.41|/\sqrt{1.09^2 + 0.32^2}$) and 1.3σ in S_8 ($|0.827 - 0.812|/\sqrt{0.010^2 + 0.006^2}$); their torsion-consistent-with-zero result parallels our ΔN_{eff} -consistent-with-zero null finding. Their AIC preference for the torsion model is *not* directly comparable to the present stack: this paper’s $\Delta\text{AIC}/\text{BIC}/\ln B$ model comparison is deliberately deferred to the nested-sampling follow-up (Sec. E3, *Model-comparison statistics* paragraph), so no model-preference statement is made here that could agree or conflict with theirs.

3. Cosmological Fits and Model Comparison

a. Datasets and Configuration

We analyze four dataset combinations: (1) Planck CMB (NPIPE (PR4) CamSpec high- ℓ TTTEEE + 2018 low- ℓ TT/EE + 2018 lensing) [6] (the NPIPE/PR4 release provides the high- ℓ CamSpec likelihood; the low- ℓ and lensing likelihoods are taken from the 2018 release, the standard pairing in the Cobaya likelihood stack); (2) +SDSS BAO — eBOSS DR16 LRG/QSO/Ly α (auto and \times QSO) [64] + SDSS DR7 MGS [65] + 6dFGS [66]; (3) +Pantheon+ [62]; (4) +SH0ES M_B -anchor [56] + DES Y3 S_8 [67]. No DESI BAO likelihood enters the frozen $\Lambda\text{CDM}+\Delta N_{\text{eff}}$ chains. The exact Cobaya likelihood blocks of all four chains are listed in Table VIII. For scope clarity: Table VII summarizes only the two *frozen* $\Lambda\text{CDM}+\Delta N_{\text{eff}}$ chains (Planck+BAO+SN and full-tension), and the Planck-only (accumulating) and Planck+BAO (diagnostic, not in headline) rows of Table VIII feed neither table. Parameter estimation uses Cobaya [68] (v3.5 original; v3.6.1 verification) with stock CAMB and ΔN_{eff} as a free parameter—no custom CAMB modifications. The extended parameter space adds ΔN_{eff} to ΛCDM , with both $(\omega/H)_0$ and Ω_k fixed to zero in the actual sampled YAML configuration (Sec. E1). Reproducibility materials at <https://github.com/Hubify-Projects/bigbounce/tree/main/reproducibility>.

b. Results: $\Lambda\text{CDM}+\Delta N_{\text{eff}}$ proxy

The two load-bearing posteriors from the frozen chains are $\Delta N_{\text{eff}} = -0.020 \pm 0.169$ (full-tension) and $+0.058 \pm 0.179$ (Planck+BAO+SN), both consistent with zero. The $\Lambda\text{CDM}+\Delta N_{\text{eff}}$ proxy thus offers neither posterior preference nor exclusion at the present data precision.

The ΔN_{eff} extension alone does not resolve the Hubble tension.

Independent re-run cross-check (this version).—A dedicated re-run of the Planck+BAO+SN $\Lambda\text{CDM}+\Delta N_{\text{eff}}$ configuration (matching likelihood stack for the SDSS BAO block and high- ℓ CamSpec component; the c15 low- ℓ EE and lensing likelihood names, `planck_2020_lollipop.lowlE` and `planckpr4lensing`, differ from the frozen-chain names in Table VIII, `planck_2018_lowl.EE` and `planck_2018_lensing.click`; the pairing mismatch is noted as a known limitation in the *Release-pairing note* below) converged on an independent compute pod to $R-1=0.0147$ (40,349 raw rows, 28,245 post-burn-in at 30%, 107,853 GetDist sum-of-weights — not the integrated-autocorrelation ESS) and gives $\Delta N_{\text{eff}} = +0.0514 \pm 0.171$, in 0.04σ agreement with the frozen $+0.058 \pm 0.179$ quote above. The auxiliary posteriors $H_0 = 67.81 \pm 1.07$ km/s/Mpc, $\sigma_8 = 0.813 \pm 0.009$, $S_8 = 0.828 \pm 0.010$, $\Omega_m = 0.311 \pm 0.006$ all reproduce the frozen-chain Table I values within $< 0.1\sigma$. The “consistent with zero / no Hubble-tension resolution” conclusion above is therefore reproduced on an independent chain; full chain artifacts and posterior summary are committed at reproducibility/cosmology/chains/wOwa_quintom_desi_dr2/c15_converged/. *Release-pairing note:* the primary frozen chains use Planck NPIPE (PR4) CamSpec high- ℓ TTTEEE paired with `planck_2018_lowl.EE` and `planck_2018_lensing.click` (Planck 2018 releases). The c15 verification rerun instead uses `planck_2020_lollipop.lowlE` and `planckpr4lensing` for the low- ℓ EE and lensing components, making c15 a release-pairing robustness rerun (PR4-consistent low- ℓ /lensing) rather than an identical-likelihood rerun of the frozen chains. The 0.04σ agreement in ΔN_{eff} across this likelihood substitution provides an empirical bound on pairing-induced bias at the quoted precision (see Sec. E2, limitation note).

Appendix F: NaMaster Pseudo- C_ℓ $E \rightarrow B$ Pipeline Validation

This appendix reports the NaMaster pseudo- C_ℓ bias-injection Monte-Carlo validation of the $E \rightarrow B$ deconvolution pipeline on synthetic ΛCDM polarization skies. It is a pipeline-recovery validation, not a competitive sky-detection significance claim.

1. Data Methods: CMB E - B Analysis

Birefringence measurements are adopted from the published literature: $\beta = 0.30^\circ \pm 0.11^\circ$ (Planck NPIPE [69]) and $\beta = 0.215^\circ \pm 0.074^\circ$ (ACT DR6 [5]). The spectator ALP analysis (Sec. G1) instead uses the Eskilt-Komatsu joint WMAP+Planck summary likelihood ($\beta =$

TABLE VIII. Per-chain Cobaya likelihood blocks, verbatim from the on-disk YAML configurations (`cobaya_*.yaml`; frozen-chain `spin_torsion.input.yaml`). All four Λ CDM+ ΔN_{eff} chains share the Planck block `planck_NPIPE_high1_CamSpec.TTTEEE + planck_2018_low1.TT + planck_2018_low1.EE + planck_2018_lensing.clik`; the SDSS BAO block is `bao.sixdf_2011_bao + bao.sdss_dr7_mgs + bao.sdss_dr16_baoplus_{lrg, qso, lyauto, lyxqso}`. The four rows are incremental: each adds to the stack of the row above.

Chain	Likelihood stack	Status
Planck-only	Planck block	accumulating (not frozen)
Planck+BAO	+ SDSS BAO block	diagnostic (not in headline)
Planck+BAO+SN	+ <code>sn.pantheonplus</code>	frozen (Table VII)
Full-tension	+ <code>H0.riess2020Mb</code> + DES-Y3 S_8 Gaussian	frozen (Table VII)

$0.342^\circ \pm 0.094^\circ$ [4])¹¹ as its primary constraint; the Planck NPIPE and ACT DR6 values above are quoted only for context and for the auxiliary inverse-variance cross-check of Eq. (G3). (The quoted 3.6σ is the ratio $\beta/\sigma = 0.342/0.094 \approx 3.6$; our likelihood is a Gaussian summary in the published mean and width, not a refit of the Eskilt–Komatsu posterior shape.)

Scope note.—*The NaMaster pseudo- C_ℓ analysis below validates only the algebraic $E \rightarrow B$ mode-coupling deconvolution on foreground-free synthetic skies; it does not establish a real-sky systematic budget (the β - α calibration degeneracy, galactic foregrounds, instrumental beams, and anisotropic noise are all absent by construction and explicitly out of scope). The high pipeline template-fit SNR figures (e.g., 20.32, 25.71; fn. 12) refer to recovery of injected MC signals and must not be conflated with the published Planck/ACT DR6 2.7–2.9 σ sky detection. The β - α degeneracy that dominates real-sky birefringence measurements is broken only by unrotated galactic foregrounds, which are absent from these synthetic skies by construction (cf. abstract).*

Pipeline configuration.—The pseudo- C_ℓ analysis follows the NaMaster framework [70]. *Simulated skies.*—Each Monte Carlo realization is a synthetic Λ CDM CMB Q/U sky generated directly at $N_{\text{side}} = 512$ (`healpy.synfast`, harmonic content to $\ell_{\text{max}} = 2N_{\text{side}} = 1024$) from a semi-analytic fit to the Planck-2018 EE spectrum plus a lensing-like BB component ($C_\ell^{BB} = 0.05 C_\ell^{EE}$). No real Planck map enters the Monte Carlo, and no instrumental beam is applied (the synthetic skies and the recovery template share the same spectra, so a common beam would *largely cancel* in the β estimate

provided the beam-deconvolution on the E -map and the C_ℓ^{EE} template are identical; a dedicated beam-mismatch MC test (varying the deconvolved-beam FWHM between map and template) is deferred to a future sky-measurement-level analysis; the same cancellation argument covers the $N_{\text{side}} = 512$ HEALPix pixel-window smoothing, which is common to the simulated maps and the spectra entering the template fit — the decoupled spectra are not pixel-window-deconvolved, and no pixel-window mismatch enters the β estimate).

Mask.—An ACT-like footprint (Galactic cut $|b| > 20^\circ$ plus declination cut $\text{dec} \in [-65^\circ, +25^\circ]$), apodized by Gaussian smoothing of the binary mask at 2° FWHM, giving $f_{\text{sky}} = 0.32$. No E/B purification is applied (`purify_b=False`); residual $E \rightarrow B$ leakage from the apodized mask is absorbed into the measured pipeline-recovery bias, which is the quantity this validation is designed to characterize.

Mode-coupling matrix and binning.—The $M_{\ell\ell'}$ matrix is computed via `NmtWorkspace.compute_coupling_matrix` on the apodized mask, retaining the full $EE/EB/BB$ block structure. (In the committed driver the Q/U maps enter as a single spin-2 `NmtField` on the apodized mask with `purify_b=False`, no beam, and all other `NmtField/NmtWorkspace` options at NaMaster library defaults — in particular no `n_iter` override; binning uses `NmtBin.from_edges` with 20 linear integer-edge bins from `np.linspace(30, 1536, 21)`; `reproducibility/p1_namaster_500mc/scripts/namaster_500mc.py`.) Spectra are band-power-binned into 20 linear bins spanning $30 \leq \ell \leq 3N_{\text{side}} = 1536$ (bins above the map band limit $\ell = 1024$ carry noise only). Only this single binning/ ℓ -range configuration is exercised; an ℓ -range robustness sweep is not part of the present MC suite.

Noise model and injections.—Each realization adds white noise at the ACT-like level $\Delta_P = 10 \mu\text{K} \cdot \text{arcmin}$ (a conservative worst-case bias check; no $1/f$ or anisotropic component). The driver script converts Δ_P to a per-pixel noise RMS as $\sigma_{\text{pix}} = \Delta_P / \sqrt{\Omega_{\text{pix}}}$ with Ω_{pix} the $N_{\text{side}} = 512$ pixel area expressed in arcmin^2 ($\Omega_{\text{pix}} = 47.21 \text{ arcmin}^2$, giving $\sigma_{\text{pix}} = 10 / \sqrt{47.21} = 1.455 \mu\text{K}$; algebraically identical to the standard $\sigma_{\text{pix}} = \Delta_P [\pi / (180 \times 60)] / \sqrt{\Omega_{\text{pix}}^{(\text{sr})}}$; Δ_P is in thermodynamic CMB units, $\mu\text{K}_{\text{CMB}} \cdot \text{arcmin}$, matching the `synfast` μK_{CMB}

¹¹ Eskilt & Komatsu 2022 disambiguation: the published PRD paper [4] (PRD 106:063503, arXiv:2205.13962) analyzes *Planck PR3 + WMAP9*; the public reproduction code released by the authors at github.com/LilleJohs/Cosmic.Birefringence was subsequently updated to use *Planck PR4 / NPIPE*. Throughout this paper, the labels “PR4/NPIPE” attached to the Eskilt+Komatsu likelihoods refer to the code-repository dataset; the ALP MCMC uses only the scalar Gaussian summary likelihood $\beta = 0.342^\circ \pm 0.094^\circ$ and does not depend on the PR3/PR4 map-level distinction except through the provenance of that summary value. The abstract $\beta = 0.342^\circ \pm 0.094^\circ$ (3.6σ) headline is from the published PR3+WMAP9 joint analysis. The repository README is the authoritative source for the dataset attribution in the executed pipeline.

maps — bandpass / colour-correction conversions do not enter the synthetic pipeline), and draws independent Gaussian realizations with the *same* σ_{pix} for Q and U (no $\sqrt{2}$ factor; `reproducibility/p1_namaster_500mc/scripts/namaster_500mc.py`). The $\beta = 0.27^\circ$, $\beta = 0.342^\circ$, and $\beta = 0$ injections rotate $Q + iU$ via $e^{2i\beta}(Q + iU)$. The angle is recovered per configuration by an unweighted χ^2 template fit of the decoupled C_ℓ^{EB} bandpowers to $\sin(2\beta)\cos(2\beta)C_\ell^{EE}$ on a β grid (a uniform grid of angles spanning $\beta \in [-2^\circ, +2^\circ]$ at 0.001° step, with β converted to radians when evaluating sin and cos; `namaster\500mc.py`). Explicitly, the estimator minimises

$$\chi^2(\beta) = \sum_b [C_b^{EB, \text{decoupled}} - \frac{1}{2}\sin(4\beta)C_b^{EE, \text{tmpl}}]^2, \quad (\text{F1})$$

where $C_b^{EE, \text{tmpl}}$ is the Planck-2018 semi-analytic EE template and the sum runs over all 20 bins $30 \leq \ell \leq 1536$. The amplitude factor $\frac{1}{2}\sin(4\beta)$ is identical to the $\sin(2\beta)\cos(2\beta)$ form quoted above (double-angle identity), and the explicit factor of $\frac{1}{2}$ is the standard cosmic-rotation EB normalization $C_\ell^{EB} = \frac{1}{2}\sin(4\beta)(C_\ell^{EE} - C_\ell^{BB})$ in the $C_\ell^{BB} \rightarrow 0$ template limit; it is present in both the equation and the analysis code and is not omitted. This form matches the canonical script exactly (`namaster_500mc.py` L223: `chi2[j] = np.sum((c1_eb_measured - c1_theory) ** 2)`) — no σ_b^2 divisor. This unweighted form is adopted to match the estimator configuration used in the public NaMaster birefringence driver scripts (e.g., [4]), facilitating direct pipeline-bias comparability; the inverse-variance-weighted variant is evaluated in the robustness battery below as a cross-check. Three clarifications on the template-band treatment: (i) the fit is *unweighted* — all bins carry equal weight regardless of their noise level (see “Canonical estimator choice” below); (ii) the template $C_b^{EE, \text{tmpl}}$ is evaluated at the same $N_{\text{side}} = 512$ pixel resolution as the simulated maps, so the $N_{\text{side}} = 512$ HEALPix pixel-window smoothing cancels identically between the decoupled $C_b^{EB, \text{decoupled}}$ and the template — no explicit pixel-window deconvolution is applied and no pixel-window mismatch enters the β estimate; (iii) bins above the map band limit $\ell_{\text{max}} = 2N_{\text{side}} = 1024$ carry zero template weight ($C_b^{EE, \text{tmpl}} = 0$ for b above this range), so the 20-bin sum is effectively restricted to $\ell \leq 1024$ despite the formal binning edge at $\ell = 1536$ — a robustness configuration explicitly confirming this (“Restricting the fit to bins with $\ell \leq 1024$ changes nothing”) is documented in the battery below. The production suite used non-negative injections; a dedicated 500-realization $\beta = -0.27^\circ$ rerun at $f_{\text{sky}} = 0.32$ (`reproducibility/p1_namaster_500mc/results/c9f_negative_beta.json`) recovers -0.238° , confirming sign-symmetric recovery with bias magnitude identical to the $+0.27^\circ$ case.

Reproducibility.—The canonical driver script, deterministic seeds (`seed.base=42`), and output summary are archived at `reproducibility/`

`p1_namaster_500mc/` (mirroring the executed pod run `pipelines/h200_results/pod1_namaster_umap_2026-04-29/`).

Production 500-realization run (April 2026).—We performed the NaMaster pseudo- C_ℓ Monte Carlo described above with 500 realizations per injection. Injecting the spectator-ALP fiducial $\beta = 0.27^\circ$ (consistent with, but not derived from, the ECH action) recovers:

$$\hat{\beta}_{\text{NaMaster}} = 0.238^\circ \quad (500\text{-MC sample mean of } \hat{\beta}). \quad (\text{F2})$$

This MC-recovery value 0.238° is a pipeline-validation figure on synthetic skies and is *not directly comparable* to the WMAP+Planck 3.6σ sky-detection significance of $\beta_{\text{obs}} = 0.342^\circ \pm 0.094^\circ$ [4]: the significance figure is a sky-measurement null-rejection statistic, whereas 0.238° is the converged MC sample mean for an *injected* signal on noise-only synthetic skies. The pipeline-recovery bias, defined throughout as $\Delta\hat{\beta} \equiv \hat{\beta} - \beta_{\text{inj}}$, is in absolute value $\leq 0.040^\circ$ across the three injection points (0.000° at $\beta_{\text{inj}} = 0$, -0.032° at 0.27° , -0.040° at 0.342°); we carry the worst case $|\Delta\hat{\beta}| = 0.040^\circ \pm 0.002^\circ$ (SE of 500-MC mean) forward as the NaMaster systematic floor (deconvolution-algebra bias on foreground-free skies; not a real-sky bias bound). This figure applies exclusively to the CMB-only, unrotated synthetic validation and supplies *no* systematic floor for real-sky β measurements, which additionally rely on galactic foregrounds to break the β - α (birefringence-miscalibration) degeneracy and require full foreground cleaning absent from these synthetic skies. At the canonical $\beta_{\text{inj}} = 0.27^\circ$ injection the bias is -0.032° (attributed by the robustness battery below to the unweighted template fit as the dominant driver, with a secondary contribution from the injected- BB realization shape — the proxy $C_\ell^{BB} = 0.05C_\ell^{EE}$ versus a lensed- Λ CDM BB spectrum — rather than a fit-template $-C_\ell^{BB}$ mismatch (carrying a $-C_\ell^{BB}$ term in the *template* produces no further shift; see the battery below), and explicitly *insensitive* to the apodization scale).¹² For $\beta = 0.342^\circ$ (the published joint

¹² The “pipeline-recovery SNR= 20.32” figure (and analogously 25.71 for the $\beta = 0.342^\circ$ injection below) is the *template-fit* SNR of the driver script, $\text{SNR}^{\text{tmpl}} \equiv [\sum_b (C_b^{EB, \text{th}}/\sigma_b^{\text{MC}})^2]^{1/2}$, where $C_b^{EB, \text{th}} = \sin(2\beta)\cos(2\beta)C_b^{EE}$ is the injected EB bandpower template and σ_b^{MC} is the per-bin scatter across the 500 realizations — the matched-template significance of the injected signal against single-realization noise, the appropriate quantity for evaluating the deconvolution pipeline. It is *not* the significance of the recovered angle itself. The sky-fraction sweep artifact (`c1_fsky_sweep.json`) confirms the identification: its `snr.template.canonical.def` values 32.98 and 28.81 at $f_{\text{sky}} = 0.85$ and 0.65 are consistent with $20.32\sqrt{f_{\text{sky}}}/0.32$ (33.12 and 28.96; within 0.5%). The *per-realization angle-recovery* ratio $\hat{\beta}/\sigma_\beta$, with σ_β the per-realization scatter of the recovered angle, is measured in the same artifact as 8.1 ($\sigma_\beta = 0.029^\circ$, $f_{\text{sky}} = 0.85$) and 7.2 ($\sigma_\beta = 0.033^\circ$, $f_{\text{sky}} = 0.65$); σ_β was not recorded for the original canoni-

WMAP+Planck value [4]), the pipeline recovers 0.302° at template-fit SNR = 25.71; for $\beta = 0$, the recovered angle is 0.000° with template-fit SNR 0.0 (null check; `summary.json`). The pipeline-recovery bias is $\Delta\hat{\beta} = -0.032^\circ$ at injection $\beta = 0.27^\circ$ ($\hat{\beta} = 0.238^\circ$) and $\Delta\hat{\beta} = -0.040^\circ$ at injection $\beta = 0.342^\circ$ ($\hat{\beta} = 0.302^\circ$). The *absolute* bias grows mildly ($\sim 25\%$) with injected amplitude, while the *multiplicative* under-recovery is amplitude-independent at $\sim 12\%$ ($0.238/0.27 = 0.302/0.342 = 0.88$). The pipeline therefore has a $\sim 12\%$ multiplicative under-recovery (0.032° absolute bias for the canonical $\beta = 0.27^\circ$ injection) with a worst-case empirical bias of $|\Delta\hat{\beta}| = 0.040^\circ$ for the unweighted estimator at the higher injection angle; we carry this 0.040° forward as the observed NaMaster pipeline bias (deconvolution-algebra bias on foreground-free skies; not a real-sky bias bound). This is a methodology cross-check, not a competitive sky measurement. (The estimator is *not* unbiased in the standard statistical sense; the 0.040° is the observed pipeline bias on the multiplicative bias, not a bound on random scatter.)

Sky-fraction sweep.—Because published birefringence analyses use larger sky fractions than our canonical $f_{\text{sky}} = 0.32$ validation mask, we repeat the $\beta = 0.27^\circ$ injection-recovery exercise at $f_{\text{sky}} \approx 0.850$ (Planck-like Galactic-cut mask: $|b| > 5^\circ$, all declinations, apodized at 2° FWHM, giving $f_{\text{sky}} \approx 0.850$; the Planck 2018 likelihood mask achieves approximately this fraction with a similar latitude threshold plus point-source excisions not included here) and $f_{\text{sky}} \approx 0.650$ (ACT-DR6-like Galactic-cut mask: $|b| > 15^\circ$, declination cut $\text{dec} \in [-60^\circ, +25^\circ]$ by design — this differs from the canonical $[-65^\circ, +25^\circ]$ mask used in the primary validation, which matches the ACT pipeline footprint; the -60° lower edge approximates the ACT DR6 survey boundary in the sky-fraction sweep; apodized at 2° FWHM) galactic-cut masks with the same apodization recipe, noise level, and $N = 500$ MC realizations ($N_{\text{side}} = 512$). The recovered values are $\hat{\beta} = 0.237^\circ$ ($f_{\text{sky}} = 0.85$, per-realization $\sigma_\beta = 0.029^\circ$) and $\hat{\beta} = 0.236^\circ$ ($f_{\text{sky}} = 0.65$, $\sigma_\beta = 0.033^\circ$), i.e. a recovery bias of -0.033° to -0.034° — statistically indistinguishable from the canonical-mask bias of -0.032° (the 0.001° and 0.002° differences are $0.8\times$ and $\approx 1.4\times$ the respective standard errors of the 500-realization means, $\sigma_\beta/\sqrt{500} \approx 0.0013^\circ$ and 0.0015°). The $\sim 12\%$ multiplicative under-recovery is therefore a

sky-fraction-independent property of the pipeline, not an artifact of the small canonical mask.

Robustness battery and bias attribution.—A six-configuration robustness battery ($N = 500$ MC each, identical seeds to the canonical run; artifact `reproducibility/p1_namaster_500mc/results/c10_robustness_battery.json`) pins down the origin of the bias. First, an independent local rerun of the canonical configuration reproduces the pod anchor exactly ($\hat{\beta} = 0.238^\circ$, bias -0.032°). The bias is then *unchanged* under an apodization-scale sweep (0.5° and 3° FWHM vs. the canonical 2° : $\hat{\beta} = 0.239^\circ$ and 0.238°), under a larger Galactic cut ($|b| > 30^\circ$, $f_{\text{sky}} = 0.20$: 0.238°), and under B -mode purification (`purify_b=True`: 0.238°) — so the earlier attribution of the residual to apodization-induced power suppression is *not* supported by direct variation of the apodization scale, and the bias is likewise independent of footprint geometry and purification. Two mechanisms do move it: (i) replacing the unweighted χ^2 template fit with an inverse-variance-weighted fit recovers $\hat{\beta} = 0.264^\circ$ (bias -0.006°), removing $\approx 80\%$ of the bias — the unweighted fit’s equal weighting of noise-dominated high- ℓ bins is the dominant contribution; and (ii) replacing the crude $C_\ell^{BB} = 0.05 C_\ell^{EE}$ proxy with a CAMB lensed- Λ CDM BB spectrum recovers $\hat{\beta} = 0.251^\circ$ (bias -0.019°), consistent with the empirical $-C_\ell^{BB}$ template-mismatch robustness test described above (the ≈ 5 percentage-point reduction in bias is the empirical effect of injecting CAMB lensed- Λ CDM BB versus the proxy $0.05 C_\ell^{EE}$, not a closed-form analytic derivation). (Clarification of what this swap varies: per the committed battery script `reproducibility/p1_namaster_500mc/scripts/c10_robustness_battery.py`, the fit template is $\sin(2\beta)\cos(2\beta)C_\ell^{EE}$ with no BB term in *all* configurations; the `camb_lensed` swap changes the BB shape of the *injected synthetic skies* only, so the test isolates the dependence of the bias on the injected- BB shape against a fixed EE -only template, not a template-shape substitution.) A matched extension config (same 500 MC skies and seeds, CAMB lensed- Λ CDM BB injected *and* a $-C_\ell^{BB}$ -bearing template $\sin(2\beta)\cos(2\beta)(C_\ell^{EE} - C_\ell^{BB})$) recovers $\hat{\beta} = 0.251^\circ$ (bias -0.019°), unchanged from the EE -only-template configuration at the 10^{-3} -degree fit-grid resolution — carrying the BB term in the template produces no further shift because the lensed BB is negligible against the synthetic EE template amplitude, so the residual -0.019° is not attributable to the remaining template-shape mismatch (config `camb_bb_template` in the same artifact). Restricting the fit to bins with $\ell \leq 1024$ changes nothing (0.238°): the noise-only bins above the band limit carry zero template weight. Because the validation carries the bias forward *empirically* as a multiplicative correction and a 0.040° observed pipeline bias (deconvolution-algebra bias on foreground-free skies; not a real-sky bias bound), none of the quoted headline numbers change; the battery

cal $f_{\text{sky}} = 0.32$ run; a dedicated 500-MC rerun at $f_{\text{sky}} = 0.32$ measures it directly as $\sigma_\beta = 0.046^\circ$ with $|\hat{\beta}|/\sigma_\beta = 5.2$ per realization (`reproducibility/p1_namaster_500mc/results/c9f_negative_beta.json`), confirming the $\sqrt{f_{\text{sky}}}$ -scaling estimate ($\approx 0.047^\circ$, ≈ 5). The pipeline-recovery SNR (template-fit significance of the injected signal) is 20.32 for the $\beta = 0.27^\circ$ injection. The same rerun injects $\beta = -0.27^\circ$ and recovers -0.238° (bias $+0.032^\circ$): recovery is sign-symmetric, with bias magnitude identical to the $+0.27^\circ$ injection. The standard error of the 500-MC mean is smaller by $\sqrt{N} = 22.4$.

sharpen only the attribution (estimator weighting + BB template shape, not apodization or masking). The total bias remains well below the published observational uncertainty ($\sigma_\beta^{\text{obs}} = 0.094^\circ$ from WMAP+Planck [4]; note $0.040^\circ \approx 0.5\sigma$ of the ACT DR6 $\sigma_\beta = 0.074^\circ$ [5], which is a published sky measurement and likewise not directly comparable to this synthetic-sky pipeline figure) at every sky fraction tested. *Canonical estimator choice.*—The unweighted χ^2 template fit is adopted as the canonical baseline to match the estimator configuration used in the public NaMaster driver scripts released by the published birefringence analyses (e.g., [4]), facilitating direct comparison of bias properties. The inverse-variance-weighted fit is evaluated in the robustness battery as a cross-check confirming that the dominant source of the $\sim 12\%$ multiplicative bias is the equal-weighting of noise-dominated high- ℓ bins; it is not adopted as the canonical estimator here precisely because that change would break direct comparability with the published estimator configuration. Artifacts: `reproducibility/p1_namaster_500mc/results/c1_fsky_sweep.json`; `reproducibility/p1_namaster_500mc/results/c10_robustness_battery.json`.

Appendix G: Spectator-ALP Cosmic-Birefringence Consistency Check

This appendix reports the spectator-ALP birefringence consistency check, including the prior-predictive accommodation-cost Monte-Carlo. The same birefringence arises in standard GR with an identical ALP; it is not a distinctive ECH prediction.

1. Cosmic Birefringence: Spectator ALP Consistency Check

Framing (prior-volume accommodation, not an independent test).—**The accommodation fraction reported here is a prior-predictive / prior-sensitivity exercise at fixed $C_{a\gamma}$, not a measurement and not a distinctive ECH prediction.** This section is an *accommodation / prior-volume* exercise, not an independent confirmation of the birefringence signal. The ALP likelihood is a Gaussian summary centred on the single published datum $\beta_{\text{obs}} = 0.342^\circ \pm 0.094^\circ$ [4], so the posterior β values ($\beta_{\text{ALP}}, \beta_{\text{free}}$) agreeing with β_{obs} “within 1σ ” is expected *by construction* and carries no independent confirmatory weight: all three numbers are constrained by the same measurement. The scientifically load-bearing outputs are therefore not the agreement itself but (i) *where* in parameter space the accommodation lives — a non-minimal photon coupling $C_{a\gamma} \gtrsim 8$ (above the standard KSVZ/DFSZ $\mathcal{O}(1)$ benchmark) and a posterior median $m \simeq 36 H_0$ well above the scan-prior midpoint — and (ii) the *prior cost* of spectator safety: only 13% of the posterior mass satisfies $\Omega_a < 0.01$, requir-

ing a $\gtrsim 100\times$ misalignment fine-tuning under a $\cos\theta_i$ -flat prior. Read this way, the result is that the model *can* accommodate the signal but only off the minimal, natural region of its own parameter space, and that the same birefringence arises identically in standard GR — it is not a distinctive ECH prediction.

Note (spectator-status caveat, main text).—The headline $\beta \approx 0.27^\circ$ benchmark of this subsection requires a $\gtrsim 100\times$ fine-tuning of the misalignment initial condition under a $\cos\theta_i$ -flat prior (equivalently $\sim 25\times$ relative to the ad-hoc $\theta_i \approx 0.5$ midpoint; $\theta_i \sim 0.1$ versus the natural-prior midpoint $\theta_i \sim 0.5$; quantitative derivation in fn. 14) to keep the ALP an actual spectator field; this caveat is load-bearing for the “consistency-not-prediction” framing and is stated here, in the main text, rather than in a footnote only. The operational criterion for spectator status throughout this section is $\Omega_a < 0.01$: samples satisfying this cut constitute 13% of the full ALP posterior mass (Table IX) and are the only subset for which the ALP dark-energy fraction is safely sub-dominant. The threshold value 0.01 is physically motivated rather than tuned to a target: $\Omega_a = \rho_a(z=0)/\rho_{\text{crit},0}$ is the present-day ALP fraction of the critical density (defined in §G 1), so $\Omega_a < 0.01$ requires the field to carry less than 1% of the energy budget today — a conservative spectator bound one full decade stricter than the nominal sub-dominance scale $\Omega_a < 0.1$. We deliberately quote the spectator-safe verdict at this stricter end. The cut-dependence is reported explicitly, not hidden: the companion subset table (Table IX) lists both the $\Omega_a < 0.1$ (44% mass, ESS = 1989) and $\Omega_a < 0.01$ (13% mass, ESS = 461) read-outs, and the rotation marginal β is stable across them (the verdict does not flip between the two thresholds); we do not attempt a Bayes-factor cut, which would require a dedicated model-comparison analysis outside the scope of this consistency check. This subsection presents a secondary consistency check using the spectator ALP model. The ALP produces cosmic birefringence independently of the gravitational theory; the same prediction $\beta \approx 0.27^\circ$ arises in standard GR with an identical ALP Lagrangian within the scan-prior envelope but near its upper-displacement/coupling edge; the posterior-supported fixed- $C_{a\gamma} = 8$ fit shifts to $m \gg H_0$ (median $\simeq 36 H_0$), with the same $\sim 25\times$ misalignment tuning required to recover the birefringence signal disclosed below and in fn. 14. It is not derived from minimal ECH (which does not produce the required photon-torsion coupling) and is not a distinctive ECH prediction. The model class was previously studied by Fujita *et al.* [71].

Headline observational constraint.—The primary observational reference adopted in this analysis is the published Eskilt & Komatsu joint WMAP+Planck value $\beta = 0.342^\circ \pm 0.094^\circ$ (3.6σ) [4] (the joint WMAP9 + Planck PR3 analysis as published; the PR4/NPIPE label refers to the repository rerun discussed in the disambiguation footnote earlier in this paper; ACT DR6 enters only via the separate $\beta = 0.215^\circ \pm 0.074^\circ$ measurement [5], used below only in the auxiliary inverse-variance

combination), which accounts for shared calibration systematics. The simplified inverse-variance combination below (3.9σ) is retained as an auxiliary cross-check only and is explicitly *not* used as the headline number anywhere in this paper.

ALP field evolution.—Numerical integration of the ALP equation of motion $\ddot{\phi} + 3H\dot{\phi} + m^2 f_a \sin(\phi/f_a) = 0$ in a Λ CDM background¹³ yields the field displacement from recombination to today, at the representative corner $m = 2H_0$, $\theta_i = 1$ of the natural-parameter box:

$$\Delta\phi/f_a \approx 0.42 \quad (m = 2H_0, \theta_i = 1). \quad (\text{G1})$$

Across the natural parameter range $m/H_0 \in [1, 3]$, $\theta_i \in [0.5, 2]$ the committed EOM grid gives $\Delta\phi/f_a \in [0.064, 1.19]$ (`research/branch_R_alp_birefringence/phase2_mcmc/c10b_alp_envelope_scan.json`).¹⁴

Birefringence value.—For $C_{a\gamma} = 8$, $\theta_i = 1$, $m \approx 3.9H_0$ (the committed EOM integration gives $\Delta\phi/f_a = 1.06$ there; artifact `research/branch_R_alp_birefringence/phase2_mcmc/c10b_alp_envelope_scan.json`; an independent fixed-step fourth-order Runge–Kutta re-integration of the same EOM in $\ln a$ with the frozen initial condition $\dot{\theta} = 0$ at $z_{\text{init}} = 3000$ reproduces $\Delta\phi/f_a = 1.0601$ to four significant figures):

$$\begin{aligned} \beta &\approx \frac{\alpha_{\text{EM}}}{4\pi} C_{a\gamma} \frac{\Delta\phi}{f_a} \\ &= \underbrace{(7.297 \times 10^{-3}/(4\pi))}_{\alpha_{\text{EM}}/(4\pi) \approx 5.81 \times 10^{-4}} \times 8 \times 1.06 \\ &= 4.93 \times 10^{-3} \text{ rad} \times \frac{180^\circ}{\pi} \approx 0.28^\circ. \end{aligned} \quad (\text{G2})$$

The product of the three factors is $5.81 \times 10^{-4} \times 8 \times 1.06 = 4.93 \times 10^{-3} \text{ rad}$ (with $\alpha_{\text{EM}} = 1/137.036$; rounding to two significant figures on $\Delta\phi/f_a$ would give

4.65×10^{-3} , so the third significant figure is sensitive to the precision of the EOM integration — here we use the four-figure $\Delta\phi/f_a = 1.0601$ from the committed Runge–Kutta re-integration). The $\alpha_{\text{EM}}/(4\pi)$ prefactor is convention-dependent ($\alpha_{\text{EM}}/(2\pi)$ appears in some Lagrangian normalizations); here it corresponds to the normalization $\mathcal{L} \supset -(g_{a\gamma}/4)\phi F_{\mu\nu}\tilde{F}^{\mu\nu}$ with $g_{a\gamma} = C_{a\gamma}\alpha_{\text{EM}}/(2\pi f_a)$ and $\beta = (g_{a\gamma}/2)\Delta\phi$, the convention fixed in the committed model specification and implemented identically throughout the numerical pipeline (`research/branch_R_alp_birefringence/02_birefringence_prediction.md`; `research/branch_R_alp_birefringence/phase2_mcmc/alp_ode.py`). The fiducial value $\beta \approx 0.27^\circ$ corresponds at $C_{a\gamma} = 8$ to EOM trajectories with, e.g., $(\theta_i = 1, m \approx 3.9H_0)$ or $(\theta_i \approx 1.4, m \approx 3H_0)$. The committed EOM integration (`research/branch_R_alp_birefringence/phase2_mcmc/alp_ode.py`) gives $\Delta\phi/f_a = 0.35\text{--}0.42$ at masses $m \approx 1.8\text{--}2H_0$ ($\theta_i = 1$); all mass pairings are verified against the released grid scan. Across the box $C_{a\gamma} \in [4, 12]$, $m/H_0 \in [1, 3]$, $\theta_i \in [0.5, 2]$ the committed EOM gives $\beta \approx 0.01\text{--}0.48^\circ$ (grid scan over physical trajectories; same artifact; the span is the *union* over the full $C_{a\gamma} \times (m/H_0, \theta_i)$ box — the lower edge is reached at $C_{a\gamma} = 4$ and the upper at $C_{a\gamma} = 12$, not at any fixed coupling), a span that brackets the observed value but reaches it only in the upper-right corner of the box ($\beta(C_{a\gamma}=8)$ peaks at 0.32° at $\theta_i = 2$, $m = 3H_0$; the observed 0.342° at $C_{a\gamma} = 8$ requires $m \gtrsim 4H_0$ or larger coupling). See Appendix H3 for the full ALP-MCMC sampled-parameter list, priors, and dataset-likelihood configuration. The spectator-consistent corner of this envelope ($\theta_i \sim 0.1$, per fn. 14) does require a $\sim 25\times$ tuning of the misalignment initial condition relative to the natural prior midpoint $\theta_i \sim 0.5$, so the model *accommodates* the observed signal but is not free of misalignment-prior tuning. The $[0.01, 0.48]^\circ$ span quoted above is computed point-by-point on physical EOM trajectories ($\Delta\phi/f_a$ is a derived function of $(m/H_0, \theta_i)$, not an independent variable); the committed grid scan gives $\beta_{\text{ALP}} \in [0.01, 0.48]^\circ$ across the full scan-prior box (the benchmark EOM grid, not the spectator-consistent subset).

Summary-likelihood combination (auxiliary cross-check).—Combining $\beta = 0.30^\circ \pm 0.11^\circ$ (Planck NPIPE [69]) and $\beta = 0.215^\circ \pm 0.074^\circ$ (ACT DR6 [5]) via inverse-variance weighting:

$$\beta_{\text{combined}} = 0.241^\circ \pm 0.061^\circ \quad (\text{naive } 3.9\sigma \text{ upper bound}) \quad (\text{G3})$$

Note: this 3.9σ is a deliberately-optimistic upper bound assuming zero Planck–ACT correlation; see Eskilt & Komatsu 3.6σ for the properly-correlated headline. (Auxiliary cross-check only; the 3.9σ is the significance of β_{combined} relative to the null $\beta = 0$, i.e. $0.241/0.061$, under the deliberately-optimistic zero-correlation assumption only.) This neglects shared calibration systematics, which produce positively correlated errors between the two measurements. Positively correlated errors *underestimate* the inverse-variance-combined σ , and there-

¹³ The ALP ODE is integrated on a Λ CDM late-time $H(z)$ here as an empirical background, distinct from the quintom-bounce dynamics that supply the early-universe / contracting-phase $H(z)$ in the underlying ECH cosmology [33]. The Λ CDM-background choice is conservative for the ALP-MCMC: a quintom-type late-time background shifts $H(z)$ at $z \lesssim 1$ by \sim few percent, propagating to a \lesssim few-percent systematic on $\Delta\phi/f_a$ — well below the $\sim 30\%$ prior-width envelope on θ_i and m/H_0 dominating the β prediction.

¹⁴ Backreaction disclosure: the ALP backreaction fraction scales as $\Omega_a \sim \rho_a/\rho_{\text{crit}} \sim (m^2 f_a^2/H_0^2 M_{\text{Pl}}^2)\theta_i^2$, so the abstract’s spectator-status restriction $\theta_i \ll 1$ (see the *Spectator-status caveat* in the abstract) implies $\Omega_a \ll 1$ only in the sub-natural sliver $\theta_i \sim 0.1$. The numerical scan range $\theta_i \in [0.5, 2]$ is RETAINED here for completeness of the parameter envelope, but the natural-prior-anchored spectator-consistent result sits at $\theta_i \sim 0.1$: at $\theta_i = 0.1$ vs the scan-midpoint $\theta_i = 0.5$ the backreaction is $\Omega_a(0.1)/\Omega_a(0.5) \sim 1/25$ (i.e., a $\sim 25\times$ fine-tuning of the misalignment initial condition is required to keep the ALP a true spectator). The $\Omega_a \sim 1$ regime at $\theta_i \sim 1$ is the dark-energy-ALP regime explicitly excluded from this spectator-consistency companion check.

fore *overestimate* the significance: the naive 3.9σ figure is an upper bound on the true significance, not a lower bound. The published joint analysis at 3.6σ [4], which properly accounts for the shared polarization-angle calibration systematics via a joint covariance matrix with the Tau A self-calibration nuisance parameters, is the headline; the 3.9σ figure here is retained only as an internal cross-check demonstrating that the uncorrelated-errors approximation is qualitatively consistent with the published result.

MCMC parameter estimation.—Dedicated MCMC sampling of the ALP parameter space (three committed configurations totaling 9,720 accepted samples: $C_{a\gamma} = 8$ fixed with θ_i and $\log_{10} m_a$ sampled, 2,160; $C_{a\gamma} \in [1, 30]$ additionally sampled, 6,840; model-independent β_{free} , 720; chain configurations and per-chain counts in Appendix H3) yields: $\beta_{\text{ALP}} = 0.336^\circ \pm 0.10^\circ$ ($C_{a\gamma} = 8$ fixed; the direct-sample priors are on the underlying ALP parameters (θ_i, m_a) , not on $\Delta\phi/f_a$, which is a derived quantity along each ALP trajectory. The MCMC posterior, anchored to the Gaussian summary likelihood on the published Eskilt–Komatsu joint WMAP+Planck measurement (fn. 11; not to the EB spectra themselves; the uniform-rotation observable is periodic, $\beta \equiv \beta + n \times 90^\circ$ for E/B , and the Gaussian summary ignores wrapping — harmless here because the posterior support is confined to $|\beta| \lesssim 0.7^\circ \ll 90^\circ$, so wrapped images carry negligible likelihood), settles at $\theta_i = 1.32 \pm 0.41$ and $m \sim 10\text{--}10^2 H_0$ (posterior median $m \approx 36 H_0$) — i.e. *outside* the natural envelope box $\theta_i \in [0.5, 2]$, $m/H_0 \in [1, 3]$ in mass — where the data-preferred joint product is $C_{a\gamma}(\Delta\phi/f_a) \approx 10.3$, corresponding at the fixed $C_{a\gamma} = 8$ to $\Delta\phi/f_a \approx 1.29$, $\sim 8\%$ above the box maximum $\Delta\phi/f_a = 1.19$ ($\theta_i = 2$, $m = 3H_0$; grid-scan artifact above); this is the same fine-tuning-prone accommodation regime flagged in fn. 14 and is consistent with the model accommodating but not naturally explaining the observed $\beta_{\text{obs}} = 0.342^\circ$, consistent with the model-independent fit $\beta_{\text{free}} = 0.344^\circ \pm 0.10^\circ$ (our internal model-independent MCMC fit, with β as a free parameter, to a Gaussian summary likelihood on the published Eskilt–Komatsu joint WMAP+Planck birefringence measurement $\beta_{\text{obs}} = 0.342^\circ \pm 0.094^\circ$ [4] — not a re-analysis of the EB spectra themselves; 720 accepted samples in the dedicated β_{free} configuration; full priors and likelihood details in Appendix H3; β_{free} denotes the unconstrained-amplitude fit distinct from β_{ALP} which has $C_{a\gamma} = 8$ fixed) and the observed $\beta_{\text{obs}} = 0.342^\circ \pm 0.094^\circ$. All three within 1σ (e.g. β_{ALP} vs β_{obs} : the model posterior mean 0.336° lies 0.06σ from the data central value, $\Delta = 0.006^\circ$, using the data uncertainty $\sigma_{\text{obs}} = 0.094^\circ$; these are consistency statements against the same single published measurement, not independent confirmations). The combined coupling-displacement product entering the birefringence formula is $C_{a\gamma}(\Delta\phi/f_a) \approx 10.3$ ($\beta = 0.342^\circ$ in radians is 5.97×10^{-3} , the prefactor $\alpha_{\text{EM}}/(4\pi)$ is 5.8×10^{-4} , giving $C_{a\gamma} \Delta\phi/f_a = \beta/[\alpha_{\text{EM}}/(4\pi)] \approx 10.3$);

with the field-displacement range $\Delta\phi/f_a \in [0.064, 1.19]$ from the committed EOM grid over the natural box, the required $C_{a\gamma}$ spans ≈ 8.6 (largest displacement: $\theta_i = 2$, $m = 3H_0$) up to ≈ 160 (smallest displacement: $\theta_i = 0.5$, $m = H_0$; note the continuous-prior scan below covers $C_{a\gamma} \in [4, 60]$ only — small-displacement corners with $\Delta\phi/f_a \lesssim 0.17$ requiring $C_{a\gamma} > 60$ lie outside that scan and carry negligible posterior support); at the posterior-preferred masses ($m \sim 10\text{--}10^2 H_0$, where the displacement saturates near $\Delta\phi/f_a \approx 1.2\text{--}1.3$), the required coupling is $\approx 8\text{--}10$. Even the lower end exceeds the standard KSVZ/DFSZ benchmark range, which predicts $|C_{a\gamma}| \sim \mathcal{O}(1)$; the entire required range therefore lies outside minimal ALP photon-coupling benchmarks and requires non-minimal model building. We emphasize that this coupling burden and the misalignment burden are *two independent* fine-tunings: a non-minimal photon coupling ($C_{a\gamma} \gtrsim 9$) is required regardless of θ_i , while the spectator-consistency restriction additionally fine-tunes the misalignment initial condition to $\theta_i \sim 0.1$ (fn. 14); neither tuning substitutes for the other. The lower end (~ 9) can be accommodated in extended models with modest photon-coupling enhancement (e.g. chiral-fermion-loop enhancement, clockwork constructions), while light-mass / small-displacement corners of the box demand couplings ($\gtrsim 50\text{--}160$) requiring substantial UV-completion enhancement. The signal is therefore accommodated across the considered parameter space rather than fine-tuned only at one benchmark, but the upper-coupling end is not generic. We emphasize that the spectator-consistent corner $\theta_i \sim 0.1$ (fn. 14) inflates the required photon-coupling further: at fixed $\beta = 0.342^\circ$, $\Delta\phi/f_a \propto \theta_i$ along the underdamped trajectory, so a $5\times$ reduction in θ_i from the scan midpoint demands a correspondingly higher $C_{a\gamma}$ to hold the $C_{a\gamma} \Delta\phi/f_a \approx 10.3$ product, pushing the required enhancement well above standard KSVZ/DFSZ $\mathcal{O}(1)$ benchmarks. Because the required coupling extends well beyond the $C_{a\gamma} \in [1, 30]$ prior of the original extended configuration (which truncated $\sim 28\%$ of the posterior mass above $C_{a\gamma} = 30$; the truncation fraction is computed from the `run2_extended` chain of Appendix H3), we reran the $C_{a\gamma}$ -free fit with a continuous uniform prior $C_{a\gamma} \in [4, 60]$ (4 MPI chains, $\hat{R} - 1 = 0.0095$, 8,955 accepted samples); this prior covers the required coupling for every trajectory with $\Delta\phi/f_a \gtrsim 0.17$, including the entire posterior-preferred saturated-displacement regime, while the small-displacement corners requiring $C_{a\gamma} > 60$ carry negligible joint-posterior support (only 5% of the posterior mass sits above $C_{a\gamma} = 55$, with no pile-up at the prior edge). The resulting posterior is broad, as expected for a single-amplitude constraint on a three-parameter degeneracy: median $C_{a\gamma} = 20.7$ with 16–84% range [7.3, 45.6]; 69% of the posterior mass falls inside [9, 51], and 22% lies below $C_{a\gamma} = 9$, consistent with the high- $\Delta\phi/f_a$ tail, where the fixed product is met at smaller coupling. These mass fractions are prior-dependent statements (flat $C_{a\gamma} \in [4, 60]$,

$\theta_i \in [0.01, \pi]$, $\log_{10} m_a$ priors), not prior-independent measurements. In particular the θ_i prior is flat in the angle; the vacuum-manifold-uniform alternative (flat in $\cos\theta_i$, the density $\propto \sin\theta_i$) carries *less* mass at small θ_i , so the quoted spectator-sliver posterior fractions would decrease further under that prior swap — the flat- θ_i choice is the more generous one for the spectator corner, and the accommodation-not-natural conclusion is robust in direction to this prior choice. A direct rerun under the $\cos\theta_i$ -flat prior (density $\propto \sin\theta_i$; 10,800 accepted samples, $\hat{R} - 1 = 0.008$) confirms this empirically: the coupling posterior is essentially unchanged (median $C_{a\gamma} = 17.1$, 16–84% [6.8, 43.4]; $\beta = 0.328^\circ \pm 0.100^\circ$) while the $\theta_i \leq 0.1$ spectator-sliver mass drops from 0.33% to 0.068% (`research/branch_R_alp_birefringence/phase2_mcmc/chains/c14_costheta/c14_summary.json`). The recovered $\beta = 0.326^\circ \pm 0.099^\circ$ posterior matches the observed $0.342^\circ \pm 0.094^\circ$, confirming the consistency-check verdict over the posterior-supported coupling range rather than at fixed benchmarks. *Spectator-subset readout* (same chain, no additional sampling): the chain’s derived ALP energy fraction Ω_a (defined in the subsection below; computed under the small-angle quadratic approximation, valid for $\theta_i \ll 1$; the leading anharmonic correction is $\mathcal{O}(\theta_i^2/12)$, i.e. $\lesssim 8\%$ at $\theta_i \sim 1$ and $\lesssim 1\%$ over the posterior-supported $\Omega_a \leq 0.01$ spectator-safe subset ($\theta_i \lesssim 0.3$; Table IX), sub-dominant to the prior width; the $\Omega_a < 0.01$ cut is applied post-sampling under this approximation, so this $\lesssim 1\%$ figure bounds the only residual model-dependence of Table IX, which a per-sample EOM-integrated ρ_a re-derivation — already validated against the approximation for the committed chain parameters — would remove entirely) satisfies $\Omega_a < 0.1$ for 44% and $\Omega_a < 0.01$ for 13% of the posterior mass; restricted to the $\Omega_a \leq 0.01$ spectator-safe subset, the rotation marginal is $\beta = 0.28^\circ \pm 0.10^\circ$ (this is the subset posterior median; the injection fiducial 0.27° is the rounded value of the canonical injection point, distinct from this posterior marginal), consistent with $\beta_{\text{obs}} = 0.342^\circ \pm 0.094^\circ$ at 0.5σ . *H_0 marginalization note:* Ω_a at each MCMC step uses the Λ CDM background Hubble rate evaluated at the Cobaya posterior mean $H_0 = 67.68 \text{ km s}^{-1} \text{ Mpc}^{-1}$ (fixed, not marginalized). Marginalizing H_0 over the Planck 1σ interval shifts Ω_a by $\lesssim 3\%$ ($\Omega_a \propto H_0^{-2}$), well below the statistical uncertainty of the $\Omega_a < 0.01$ cut. The strict $\theta_i \leq 0.1$ sliver of fn. 14 carries only 0.33% of the posterior mass by MC weight (42 of the 8,955 raw samples, i.e. 0.47% by raw count — the weighted fraction is lower because the sliver samples carry below-average MC weights — too few for a stable marginal, so the following band is indicative only): within that sliver the required coupling piles against the upper prior edge, with weighted $C_{a\gamma}$ 16/50/84 percentiles of 36.8/47.2/55.6 (flat prior $C_{a\gamma} \in [4, 60]$), well above the KSVZ/DFSZ $\mathcal{O}(1)$ benchmarks (`research/branch_R_alp_birefringence/phase2_mcmc/c10a_spectator_slice.json`). This

quantifies the misalignment tuning of the spectator-consistent corner. Convergence: $\hat{R} - 1 < 0.01$ for all runs. The joint posterior structure of this continuous-prior configuration is shown in Fig. 11. Artifact: `research/branch_R_alp_birefringence/phase2_mcmc/chains/c5_continuous/`.

ALP dark-energy fraction Ω_a : definition and computation

The spectator-status classification in this section uses the ALP dark-energy fraction $\Omega_a = \rho_a(z=0)/\rho_{\text{crit},0}$. We define this quantity precisely here because the Table IX entries (44% at $\Omega_a < 0.1$; 13% at $\Omega_a < 0.01$) depend on it directly, and two independent reviewers flagged the absence of an explicit derivation.

Axion potential.—The spectator ALP carries the standard cosine potential

$$V(\phi) = m_a^2 f_a^2 \left[1 - \cos\left(\frac{\phi}{f_a}\right) \right], \quad (\text{G4})$$

where m_a is the ALP mass, f_a the decay constant, and $\phi_i = \theta_i f_a$ the initial field value parameterized by the misalignment angle $\theta_i \equiv \phi_i/f_a$. For $\theta_i \ll 1$ the potential is approximately quadratic, $V \approx \frac{1}{2} m_a^2 \phi^2$; for $\theta_i \sim \mathcal{O}(1)$ anharmonic corrections enter at $\mathcal{O}(\theta_i^2/12)$.

Onset of oscillations.—The ALP is frozen by Hubble friction and begins coherent oscillations when

$$3H(z_{\text{osc}}) = m_a, \quad (\text{G5})$$

where $H(z)$ is evaluated on the Λ CDM background (with $H_0 = 67.68 \text{ km s}^{-1} \text{ Mpc}^{-1}$, fixed at the Cobaya posterior mean; see the H_0 marginalization note in the spectator-subset readout above). For ALP masses in the scan prior $m/H_0 \in [7 \times 10^{-3}, 7 \times 10^2]$, the onset redshift ranges from $z_{\text{osc}} \lesssim 0$ (lightest masses, ALP still frozen; behaves as a cosmological-constant contribution during the survey epoch) to $z_{\text{osc}} \gg 1$ (heaviest masses, early dark matter). Equation (G5) is solved numerically at each MCMC step in `research/branch_R_alp_birefringence/phase2_mcmc/alp_ode.py`.

Energy density today.—Once oscillating ($z \ll z_{\text{osc}}$), the ALP cycle-averaged energy density redshifts as matter:

$$\rho_a(z) = \rho_a(z_{\text{osc}}) \left(\frac{1+z}{1+z_{\text{osc}}} \right)^3, \quad z \ll z_{\text{osc}}. \quad (\text{G6})$$

At onset the potential dominates: $\rho_a(z_{\text{osc}}) \approx V(\phi_i) = m_a^2 f_a^2 [1 - \cos(\theta_i)]$. The dark-energy fraction today is therefore

$$\Omega_a \equiv \frac{\rho_a(z=0)}{\rho_{\text{crit},0}} \approx \frac{m_a^2 f_a^2 [1 - \cos(\theta_i)]}{\rho_{\text{crit},0} (1+z_{\text{osc}})^3}, \quad (\text{G7})$$

with $\rho_{\text{crit},0} = 3H_0^2 M_{\text{Pl}}^2 \approx 3.7 \times 10^{-11} \text{ eV}^4$ and $M_{\text{Pl}} \approx 2.44 \times 10^{18} \text{ GeV}$. For $f_a = M_{\text{Pl}}$ and small θ_i this gives

$\Omega_a \approx m_a^2 \theta_i^2 / (6H_0^2(1 + z_{\text{osc}})^3)$, confirming the $\Omega_a \propto \theta_i^2$ scaling of fn. 14.

Computation from sampled parameters.—At each MCMC step the ALP module receives (m_a, θ_i) (with $f_a = M_{\text{Pl}}$ fixed), solves Eq. (G5) for z_{osc} , and evaluates Eq. (G7) using the potential-dominated approximation, verified against full EOM integration for committed chain parameters. For the lightest masses with $z_{\text{osc}} \leq 0$ the ALP remains Hubble-frozen through the survey epoch and is held at $\rho_a = V(\theta_i)$ (a cosmological-constant-like contribution, as noted in the onset discussion above) rather than diluted by the $(1 + z_{\text{osc}})^{-3}$ factor of Eq. (G7). The per-step Ω_a values are stored in `research/branch_R_alp_birefringence/phase2_mcmc/chains/c5_continuous/` and define the spectator-status cuts in Table IX below. The 44%/13% posterior-mass fractions are direct weighted integrals over these per-step values; they are prior-dependent (flat $\theta_i \in [0.01, \pi]$) and not prior-independent measurements of Ω_a .

LiteBIRD forecast.—LiteBIRD is projected to achieve $\sigma(\beta) \approx 0.03^\circ$ [31] (a forecast under foreground- and calibration-control assumptions, not a guaranteed instrument performance). For $\beta = 0.27^\circ$ this corresponds to a *null-rejection* significance of $\sim 9\sigma$ against $\beta = 0$, paired with only $\approx 0.7\sigma$ model-discrimination separation from the current WMAP+Planck central value ($|0.342 - 0.27|/\sqrt{0.03^2 + 0.094^2} \approx 0.7\sigma$): these two figures address *distinct* null hypotheses, and the 9σ is a detection-of-nonzero- β forecast, *not* a forecast of model discrimination between the spectator-ALP value and the published central value.

Caveats.—This birefringence prediction is independent of bounce cosmology: the ALP is a spectator field that does not participate in the bounce dynamics. The ECH framework provides heuristic motivation ($f_a \sim M_{\text{Pl}}$ from the Holst sector pseudoscalar structure) but no derived photon-torsion coupling connects the Holst action to a specific ALP potential. The model can *accommodate* the observed β within the scan-prior envelope but near its upper-displacement/coupling edge; the posterior-supported fixed- $C_{a\gamma} = 8$ fit shifts to $m \gg H_0$ (median $\simeq 36 H_0$). The spectator-safe interpretation requires a tuned misalignment subspace (the $\sim 25\times$ misalignment tuning required for the birefringence signal is disclosed in Sec. G1 and fn. 14) and a non-minimal photon coupling; it does not uniquely predict the signal.

Appendix H: Data Availability, Reproducibility Materials, and Sampled-Parameter Priors

This appendix collects the unified data-availability statement for the folded MCMC, NaMaster, and spectator-ALP analyses: the public reproducibility repository structure, the per-claim classification, and the full ALP-MCMC sampled-parameter, prior, and likelihood stack. All computational artifacts behind the num-

bers reported in Appendices E–G are archived here.

1. Reproducibility Materials

Repository structure.—The public repository <https://github.com/Hubify-Projects/bigbounce> contains:

- Four Cobaya YAML configurations (`cobaya_planck.yaml`, `cobaya_planck_bao.yaml`, `cobaya_planck_bao_sn.yaml`, `cobaya_full_tension.yaml`)—stock CAMB, no torsion modifications.
- [reproducibility/galaxy_spins/spin_fit_stan.py](#)—hierarchical Bayesian model (CmdStanPy) fitting $A(z)$ to published aggregate CW/CCW galaxy counts (program-wide content used by Paper IV, not by this paper’s analyses).
- [research/data_build/build_galaxy_spin_dataset.py](#)—reproducible pipeline downloading Galaxy Zoo DECaLS [72] from Zenodo (DOI: 10.5281/zenodo.4573248, CC-BY-4.0; likewise Paper IV content).
- [reproducibility/docs/IMPLEMENTATION_MAP.md](#)—mapping from each result to its code artifact.
- [reproducibility/docs/KNOWN_GAPS.md](#)—honest disclosure of what cannot currently be reproduced.

What is included vs. regenerable.—The two frozen $\Lambda\text{CDM} + \Delta N_{\text{eff}}$ chain directories are committed in-repo and contain the chains and diagnostics that back Table VII, Table VIII, and the relevant Table X claim-classification entries of this paper; the ALP `c5` continuous chain backs Table IX separately: `reproducibility/cosmology/frozen/full_tension_20260311_1728/` and `reproducibility/cosmology/frozen/planck_bao_sn_20260312_1954/` (each with `chains/chain_0[1-6]/` and `diagnostics/parameter_summary_CORRECTED.json`). The ALP-MCMC chains are at `research/branch_R_alp_birefringence/phase2_mcmc/chains/`. Fresh $\Lambda\text{CDM} + \Delta N_{\text{eff}}$ proxy chains for independent re-verification are NOT bundled and must be regenerated locally via `reproduce_cosmology.sh` (~ 4 – 12 h per config on 4 CPU cores). All Table VII numerical entries quoted in this paper were recomputed directly from the committed raw chains (and from the `parameter_summary_CORRECTED.json` artifacts derived from them), not from the legacy column-permuted export. No CNN galaxy classifier is included; the hierarchical fit uses published catalog labels. No CMB polarization map analysis code is provided beyond the NaMaster driver script; all published birefringence values are literature citations.

HuggingFace datasets.—The following HuggingFace datasets accompany this work (DOI assignment is pending; identifiers will be inserted at submission; no DOI

TABLE IX. Restricted-posterior readout of the continuous-prior `c5_continuous` chain on three spectator-status subsets, together with the full chain. Posterior-mass columns give the fraction of MC weight (and, for the strict $\theta_i \leq 0.1$ sliver, raw-count fraction in parentheses); β medians are direct chain readouts; m/H_0 , θ_i , and $C_{a\gamma}$ entries report medians or representative weighted statistics; ESS is the weight-expanded Sokal effective-sample size for the marker parameter on the indicated subset. *The fixed-coupling ($C_{a\gamma} = 8$) posterior-supported median sits at $m \simeq 36 H_0$ — well above the scan-prior $m \sim H_0$ benchmark — and only the $\Omega_a \leq 0.01$ spectator-safe subset enforces a true spectator while still consistent with the observed $\beta_{\text{obs}} = 0.342^\circ \pm 0.094^\circ$ at $\sim 0.5\sigma$. The spectator-consistent verdict rests on this $\Omega_a < 0.01$ subset (13% mass, stable ESS = 461), not on the strict $\theta_i \leq 0.1$ sliver (0.33% weighted mass, 42 samples), which is reported as “indicative only.” The strict-sliver row (0.33% weighted mass, 42 raw samples, $C_{a\gamma}$ percentiles 36.8/47.2/55.6) is machine-verified against the committed reduction `c10a_spectator_slice.json` (input chains `c5.[1-4].txt`, 8,955 total samples); the coarser $\Omega_a < 0.1$ and $\Omega_a < 0.01$ mass-fraction and ESS entries are illustrative weighted-integral readouts of the same continuous-prior chain and are not backed by a separate committed summary artifact. All posterior-mass fractions in this table are conditional on the Gaussian summary-likelihood approximation for the Eskilt–Komatsu β_{obs} datum (Appendix H 3) and therefore omit EB-specific band-power covariance and calibration systematics; they would shift under a full re-analysis of the underlying EB spectra.*

Subset	post. mass	β (deg)	m/H_0	θ_i	$C_{a\gamma}$	ESS
full chain	100%	0.326 ± 0.099	median $\simeq 36$	broad	med. 20.7, [7.3, 45.6]	β 2860; θ_i 796
$\Omega_a < 0.1$	44%	0.328 ± 0.100	4.7/37.7/264	0.22/0.41/0.70	14.2/26.2/46.4	1989
$\Omega_a < 0.01$ (safe)	13%	0.28 ± 0.10	6.0/40.5/238	0.15/0.21/0.27	29.9/43.3/54.1	461
$\theta_i \leq 0.1$ (strict)	0.33% (0.47% raw)	indicative only	$\sim \mathcal{O}(1)$	≤ 0.1 (tuned)	16/50/84 = 36.8/47.2/55.6	sliver-only (42 samples)

is fabricated here. In the interim, the stable references are the frozen-chain git commits and repository SHAs pinned in `CHANGELOG.md` and `reproducibility/`; these are version-locked and independently reproducible without a DOI. The URLs are also recorded in the repository `CHANGELOG.md` under the entry for `v1U.0.1` and are preserved against future README changes via the version-stamp commit identified by the `changelog`:

1. MCMC chain diagnostics and convergence CSV files: <https://huggingface.co/datasets/bamfai/p1b-mcmc-diagnostics>.
2. NaMaster pipeline artifacts (mask, MC seeds, output spectra): <https://huggingface.co/datasets/bamfai/p1b-namaster-artifacts>.
3. ALP parameter MCMC chains: <https://huggingface.co/datasets/bamfai/p1b-alp-chains>.

2. Claims Classification

Table X classifies the principal load-bearing quantitative claims made in this companion by claim type and verification status; it is the machine-checkable index used by the reproducibility audits of Appendix H 1.

3. ALP-MCMC Sampled Parameters, Priors, and Likelihood Stack

The ALP-MCMC results quoted in Sec. G 1 ($\beta_{\text{ALP}} = 0.336^\circ \pm 0.10^\circ$ at $C_{a\gamma} = 8$ fixed; $\beta_{\text{free}} = 0.344^\circ \pm 0.10^\circ$ model-independent; 9,720 total accepted samples across the three committed configurations) use the following setup. All priors below are read directly from the archived chain configurations (`research/branch_R_`

`alp_birefringence/phase2_mcmc/chains/`); the headline posteriors are verified directly against the committed chains.

Configuration (i) — fixed-coupling fit (`run1_full`; 2,160 accepted samples).

- $C_{a\gamma}$: fixed at 8 (theory-module default; not sampled).
- θ_i : uniform prior on $[0.01, \pi]$.¹⁵
- $\log_{10}(m_a/\text{eV})$: uniform prior on $[-35, -30]$ ($m/H_0 \approx 7 \times 10^{-3}$ to 7×10^2).
- f_a : fixed at M_{Pl} (spectator-class theoretical input from the Holst-sector pseudoscalar structure; not sampled).

This configuration yields $\beta_{\text{ALP}} = 0.336^\circ \pm 0.10^\circ$ with posterior $\theta_i = 1.32 \pm 0.41$ and median $m \approx 36 H_0$.

Configuration (ii) — sampled-coupling fit (`run2_extended`; 6,840 accepted samples).—As configuration (i) with $C_{a\gamma}$ additionally sampled, uniform prior $[1, 30]$. The $[4, 60]$ continuous-prior rerun (below) is the primary coupling-inference result; the $[1, 30]$ prior truncated $\sim 28\%$ of the coupling posterior mass.

Configuration (iii) — model-independent β_{free} fit (`run3_baseline`; 720 accepted samples).

- β : uniform prior on $[-1^\circ, 2^\circ]$; sampled as a free amplitude with no ALP-model parametric structure. Yields $\beta_{\text{free}} = 0.344^\circ \pm 0.10^\circ$.

Sampled parameters and priors (continuous-prior cross-check configuration; Fig. 11).—Three sampled

¹⁵ Backreaction / spectator-status disclosure: see fn. 14 in Sec. G 1. The wide prior is an envelope-completeness choice, not the spectator-consistent sub-range ($\theta_i \sim 0.1$, a $\sim 25\times$ tuning relative to the natural midpoint); posterior samples at $\theta_i \gtrsim 0.5$ belong to the dark-energy-ALP regime excluded from the spectator-consistency claim.

TABLE X. Claims classification for this appendix. The “Reference value” column identifies the public source or internal verification anchor for each claim; “Int. verified” means the quoted value has been reproduced from committed chains/artifacts at the present `v1U.0.1` commit (`b22f8cc9`); a public tagged release is pending (see Data and Code Availability).

Claim	Type	Reference value	Notes
$\Delta N_{\text{eff}} = -0.020 \pm 0.169$ (full-tension)	MCMC	Int. verified (frozen chains)	Stock CAMB proxy
$\Delta N_{\text{eff}} = +0.058 \pm 0.179$ (Planck+BAO+SN)	MCMC	Int. verified (frozen chains)	Stock CAMB proxy
$H_0 = 67.68 \pm 1.06$ (full-tension)	MCMC	Int. verified; Planck-consistent	Recovers Λ CDM
$H_0 = 67.78 \pm 1.09$ (Planck+BAO+SN)	MCMC	Int. verified; Planck-consistent	Recovers Λ CDM
Model-comparison $\Delta\text{AIC}/\text{BIC}/\ln B$	Numerical	Not reported	Nested-sampling follow-up
$\hat{\beta}_{\text{NaMaster}} = 0.238^\circ$ (500-MC)	Numerical	Int. verified (pipeline)	Obs. pipeline bias
$\beta_{\text{ALP}} = 0.336^\circ \pm 0.10^\circ$	MCMC	Int. verified (ALP chains)	ALP MCMC
Published 3.6σ ($\beta = 0.342 \pm 0.094^\circ$)	Lit.	[4]	Eskilt et al.
Stock CAMB proxy \neq ECH theory module	Scope	Definition	§E 2
ALP birefringence not distinctive ECH prediction	Scope	Definition	§G 1

parameters (read directly from the archived chain configuration, `research/branch_R_alp_birefringence/phase2_mcmc/chains/c5_continuous/c5.input.yaml`):

- $C_{a\gamma}$: uniform prior on [4, 60] (covering the posterior-supported band [9, 51] that contains 69% of the posterior mass, see Sec. G 1; the full kinematic natural-box requirement extends to $C_{a\gamma} \approx 160$ at the smallest misalignment displacements, as discussed in the same section).
- θ_i : uniform prior on [0.01, π].
- $\log_{10}(m_a/\text{eV})$: uniform prior on $[-35, -30]$ ($m/H_0 \approx 7 \times 10^{-3}$ to 7×10^2).

f_a is not a sampled parameter in this chain. The likelihood is the same Eskilt–Komatsu Gaussian summary as above; 8,955 accepted samples, $\hat{R} - 1 = 0.0095$.

Effective sample sizes (ESS).—Per-parameter ESS computed from the integrated autocorrelation time of the weight-expanded chains (Sokal estimator; `tools/ess_alp_chains.py`):

Parameter	Chain	N_{acc}	ESS
θ_i	<code>c5_continuous</code>	8,955	796
$\log_{10}(m_a/\text{eV})$	<code>c5_continuous</code>	8,955	1,114
$C_{a\gamma}$	<code>c5_continuous</code>	8,955	814
β_{deg} (derived)	<code>c5_continuous</code>	8,955	2,860
β_{free}	<code>run3_baseline</code>	720	265

The `c5_continuous` ESS values (~ 800 – 2860) are adequate for posterior characterisation. The β_{free} `run3_baseline` chain (720 accepted samples, $\text{ESS} \approx 265$) is marginal; the result $\beta_{\text{free}} = 0.344^\circ \pm 0.10^\circ$ should be interpreted with the caveat that ~ 720 accepted samples provides limited ESS for this single-parameter fit.

Likelihood stack.—All fits use a Gaussian summary likelihood on the published Eskilt–Komatsu joint WMAP+Planck isotropic-birefringence measurement $\beta_{\text{obs}} = 0.342^\circ \pm 0.094^\circ$ [4] (fn. 11; the chain configuration encodes exactly `beta_obs: 0.342, sigma_beta: 0.094`) — i.e. the constraint enters at the level of the

published β posterior, not as a direct re-analysis of the EB spectra. *Effect of the summary-likelihood approximation*.—Because a single Gaussian datum ($\beta_{\text{obs}}, \sigma_\beta$) constrains only the combination $C_{a\gamma} \Delta\phi/f_a \propto \beta$ that fixes the ALP birefringence amplitude, replacing it with the full joint EB likelihood would principally re-weight the *width* and tails of the amplitude posterior rather than translate its central location: the accommodation fractions (11.6% within 1σ at fixed $C_{a\gamma} = 8$; the $\Omega_a < 0.01$ subset fractions) therefore inherit this approximation and should be read as summary-likelihood estimates, whereas the reported posterior medians ($m \simeq 36 H_0$, $C_{a\gamma} \Delta\phi/f_a \approx 10.3$) are set by the mean of the same Eskilt–Komatsu datum and are comparatively insensitive to it. A full joint- EB re-fit to quantify the residual median shift is left as a dedicated follow-up; it is not expected to move the medians beyond their quoted 16–84% ranges. The separate ACT DR6 measurement $\beta = 0.215^\circ \pm 0.074^\circ$ [5] is an independent cross-check quoted in Sec. G 1 and does *not* enter any MCMC likelihood. The MCMC engine is Cobaya v3.6.1 with the Metropolis-Hastings sampler; convergence threshold $\hat{R} - 1 < 0.01$ across configurations (i)–(iii) ($N_{\text{tot}} = 9,720$ accepted samples) and for the continuous-prior $C_{a\gamma} \in [4, 60]$ configuration ($\hat{R} - 1 = 0.0095$, 8,955 accepted samples). *Chain-to-number map*: $\beta_{\text{ALP}} = 0.336^\circ \pm 0.10^\circ$ — the fixed- $C_{a\gamma} = 8$ chain `run1_full` (2,160 samples); $\beta_{\text{free}} = 0.344^\circ \pm 0.10^\circ$ — the model-independent single-parameter chain `run3_baseline` (720 samples); median $C_{a\gamma} = 20.7$, 16–84% range [7.3, 45.6], $\beta = 0.326^\circ \pm 0.099^\circ$, and the spectator-subset readouts — the continuous-prior `c5` chain (8,955 samples); the $C_{a\gamma} \in [1, 30]$ chain `run2_extended` (6,840 samples) is retained for the prior-truncation comparison only.

Scope statement.—These fits constrain a specific ultra-light ALP class ($m \sim H_0$, $f_a \sim M_{\text{Pl}}$; the $C_{a\gamma} \in [4, 12]$ benchmark sweep plus the continuous-prior $C_{a\gamma} \in [4, 60]$ headline configuration). The consistency with β_{obs} established here is *not* a universal mechanism-independent statement; alternative parity-violating mechanisms (Chern-Simons, axion dark matter

at higher m , etc.) require independent fits with their own model parameters.

-
- [1] Y.-F. Cai, W. Xue, R. Brandenberger, and X. Zhang, Non-gaussianity in a matter bounce, *JCAP* **0905**, 011, [arXiv:0903.0631](#).
- [2] H. Golden, $f_{NL} = -35/8$ Forecast: SPHEREx Discrimination of Bounce vs. Inflation, (2026), companion paper, posted concurrently on arXiv.
- [3] 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 \[astro-ph.CO\]](#).
- [4] 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 \[astro-ph.CO\]](#).
- [5] P. Diego-Palazuelos and E. Komatsu, Cosmic birefringence from the Atacama Cosmology Telescope data release 6, arXiv preprint (2025), [arXiv:2509.13654 \[astro-ph.CO\]](#).
- [6] Planck Collaboration, N. Aghanim, *et al.*, Planck 2018 results. VI. cosmological parameters, *Astronomy & Astrophysics* **641**, A6 (2020), [arXiv:1807.06209 \[astro-ph.CO\]](#).
- [7] S. Weinberg, The cosmological constant problem, *Reviews of Modern Physics* **61**, 1 (1989).
- [8] DESI Collaboration, A. G. Adame, *et al.*, DESI 2024 VI: cosmological constraints from the measurements of baryon acoustic oscillations, arXiv preprint (2024), [arXiv:2404.03002 \[astro-ph.CO\]](#).
- [9] DESI Collaboration, M. Abdul-Karim, *et al.*, DESI DR2 results II: Measurements of baryon acoustic oscillations and cosmological constraints, *Physical Review D* **112**, 083515 (2025), [arXiv:2503.14738 \[astro-ph.CO\]](#).
- [10] H. Golden, Big Bounce reproducibility archive: MCMC chains, NaMaster EB validation, spectator-ALP fit, and galaxy/anomaly catalogs, <https://github.com/Hubify-Projects/bigbounce/tree/main/reproducibility> (2026), committed computational artifacts underlying the companion-imported numbers of Table 3; archived with this submission.
- [11] A. Ashtekar and P. Singh, Loop quantum cosmology: A status report, *Classical and Quantum Gravity* **28**, 213001 (2011), [arXiv:1108.0893 \[gr-qc\]](#).
- [12] 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).
- [13] N. J. Popławski, Cosmological constant from quarks and torsion, *Annalen der Physik* **523**, 291 (2011), [arXiv:1005.0893 \[gr-qc\]](#).
- [14] N. J. Popławski, Universe in a black hole in Einstein-Cartan gravity, *The Astrophysical Journal* **832**, 96 (2016), [arXiv:1410.3881 \[gr-qc\]](#).
- [15] S. Mercuri, Peccei-quinn mechanism in gravity and the nature of the Barbero-Immirzi parameter, *Physical Review Letters* **103**, 081302 (2009), [arXiv:0902.2764 \[gr-qc\]](#).
- [16] 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 \[hep-th\]](#).
- [17] A. Ashtekar, J. C. Baez, A. Corichi, and K. Krasnov, Quantum geometry and black hole entropy, *Physical Review Letters* **80**, 904 (1998), [arXiv:gr-qc/9710007](#).
- [18] M. Domagała and J. Lewandowski, Black-hole entropy from quantum geometry, *Classical and Quantum Gravity* **21**, 5233 (2004), [arXiv:gr-qc/0407051](#).
- [19] K. A. Meissner, Black-hole entropy in loop quantum gravity, *Classical and Quantum Gravity* **21**, 5245 (2004), [arXiv:gr-qc/0407052](#).
- [20] 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 \[gr-qc\]](#).
- [21] D. Saadeh, S. M. Feeney, A. Pontzen, H. V. Peiris, and J. D. McEwen, How isotropic is the universe?, *Physical Review Letters* **117**, 131302 (2016), [arXiv:1605.07178 \[astro-ph.CO\]](#).
- [22] V. A. Kuzmin, V. A. Rubakov, and M. E. Shaposhnikov, On anomalous electroweak baryon-number non-conservation in the early universe, *Phys. Lett. B* **155**, 36 (1985).
- [23] H. Golden, Galaxy Chirality at Scale: 8.47M Galaxies Classified, Hemisphere Null at $p_{LEE} < 10^{-4}$, (2026), companion paper, posted concurrently on arXiv.
- [24] F. W. Hehl and B. K. Datta, Nonlinear spinor equation and asymmetric connection in general relativity, *J. Math. Phys.* **12**, 1334 (1971).
- [25] R. Jackiw and S.-Y. Pi, Chern-simons modification of general relativity, *Physical Review D* **68**, 104012 (2003).
- [26] S. Holst, Barbero's Hamiltonian derived from a generalized Hilbert-Palatini action, *Physical Review D* **53**, 5966 (1996), [arXiv:gr-qc/9511026 \[gr-qc\]](#).
- [27] G. Date, R. K. Kaul, and S. Sengupta, Topological interpretation of Barbero-Immirzi parameter, *Phys. Rev. D* **79**, 044008 (2009), [arXiv:0811.4496 \[gr-qc\]](#).
- [28] D. Benedetti and S. Speziale, Perturbative quantum gravity with the Immirzi parameter, *JHEP* **06**, 107, [arXiv:1104.4028 \[hep-th\]](#).
- [29] D. Benedetti and S. Speziale, Perturbative running of the Immirzi parameter, *J. Phys. Conf. Ser.* **360**, 012011 (2012), [arXiv:1111.0884 \[hep-th\]](#).
- [30] A. Lue, L. Wang, and M. Kamionkowski, Cosmological signature of new parity violating interactions, *Phys. Rev. Lett.* **83**, 1506 (1999), [arXiv:astro-ph/9812088 \[astro-ph\]](#).
- [31] LiteBIRD Collaboration, E. Allys, *et al.*, Probing cosmic inflation with the LiteBIRD cosmic microwave background polarization survey, *Progress of Theoretical and Experimental Physics* **2023**, 042F01 (2023), [arXiv:2202.02773 \[astro-ph.IM\]](#).
- [32] S. M. Carroll, Quintessence and the rest of the world: Suppressing long-range interactions, *Physical Review Letters* **81**, 3067 (1998), [arXiv:astro-ph/9806099 \[astro-ph\]](#).
- [33] 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 \[hep-th\]](#).
- [34] L. Shamir, Analysis of the alignment of non-random patterns of spin directions in populations of spiral galaxies,

The Astrophysical Journal **938**, 77 (2022).

- [35] L. Shamir, Asymmetry in galaxy spin directions in JWST JADES data, arXiv preprint (2024), arXiv:2401.09450 [astro-ph.GA].
- [36] D. Patel and H. Desmond, A critical assessment of galaxy spin asymmetry studies, *Monthly Notices of the Royal Astronomical Society* **528**, 2553 (2024).
- [37] O. H. E. Philcox and J. Ereira, Testing cosmic parity violation with galaxy spins, *Physical Review D* **111**, 023501 (2025), arXiv:2410.18185 [astro-ph.CO].
- [38] C. Heinrich, O. Dore, and E. Krause, Measuring f_{nl} with the spherex multi-tracer redshift space bispectrum, *JCAP* **2024** (04), 074, arXiv:2311.13082 [astro-ph.CO].
- [39] S. Dehghani, G. Geshnizjani, and J. Quintin, Cuscuton Bounce Beyond the Linear Regime: Bispectrum and Strong Coupling, (2025), arXiv:2503.01992 [gr-qc].
- [40] K. Gödel, An Example of a New Type of Cosmological Solutions of Einstein's Field Equations of Gravitation, *Rev. Mod. Phys.* **21**, 447 (1949).
- [41] N. J. Poplawski, Cosmology with torsion: An alternative to cosmic inflation, *Physics Letters B* **694**, 181 (2010), arXiv:1007.0587 [astro-ph.CO].
- [42] S. Mercuri, Fermions in the Ashtekar-Barbero connection formalism for arbitrary values of the Immirzi parameter, *Physical Review D* **73**, 084016 (2006), arXiv:gr-qc/0601013 [gr-qc].
- [43] T. Liu, X. Li, T. Xu, M. Biesiada, and J. Wang, Torsion cosmology in the light of DESI, supernovae and CMB observational constraints, *European Physical Journal C* (2025), arXiv:2507.04265 [gr-qc].
- [44] S. Legner, W. Handley, and W. Barker, Alleviating the Hubble tension with torsion condensation (TorC), arXiv e-prints (2025), arXiv:2507.09228 [astro-ph.CO].
- [45] S. Alam, S. Sen, and S. Sengupta, Bouncing cosmologies in modified gravity with space time torsion, *Eur. Phys. J. C* (2025), arXiv:2509.03508 [gr-qc].
- [46] Y.-F. Cai and J.-H. Zhu, Smoking-gun signatures of bounce cosmology from echoes of relic gravitational waves, (2026), arXiv:2603.13924 [astro-ph.CO].
- [47] T. Papanikolaou, S. Banerjee, Y.-F. Cai, S. Capozziello, and E. N. Saridakis, Primordial black holes and induced gravitational waves in non-singular matter bouncing cosmology, *JCAP* **06**, 066, arXiv:2404.03779 [gr-qc].
- [48] H. Golden, Spectrally Unusual Sources at Scale: A Multi-Survey Catalog of 378,280 Anomalies and Native-Trained Novelty Rates from 37.3 Million Sources, (2026), companion paper, posted concurrently on arXiv.
- [49] 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].
- [50] I. Brivio and M. Trott, The Standard Model as an Effective Field Theory, *Phys. Rept.* **793**, 1 (2019), arXiv:1706.08945 [hep-ph].
- [51] 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].
- [52] C. Itzykson and J.-B. Zuber, *Quantum Field Theory* (McGraw-Hill, 1980).
- [53] J. F. Nieves and P. B. Pal, Generalized fierz identities, *Am. J. Phys.* **72**, 1100 (2004), arXiv:hep-ph/0306087 [hep-ph].
- [54] S. M. Carroll, G. B. Field, and R. Jackiw, Limits on a Lorentz- and parity-violating modification of electrodynamics, *Physical Review D* **41**, 1231 (1990).
- [55] D. Harari and P. Sikivie, Effects of a nambu-goldstone boson on the polarization of radio galaxies and the cosmic microwave background, *Phys. Lett. B* **289**, 67 (1992).
- [56] A. G. Riess, W. Yuan, L. M. Macri, *et al.*, A comprehensive measurement of the local value of the Hubble constant with 1 km/s/Mpc uncertainty from the Hubble Space Telescope and the SH0ES team, *The Astrophysical Journal Letters* **934**, L7 (2022), arXiv:2112.04510 [astro-ph.CO].
- [57] N. J. Poplawski, Four-fermion interaction from torsion as dark energy, *General Relativity and Gravitation* **44**, 491 (2012), arXiv:1102.5667 [gr-qc].
- [58] M. P. Unger and N. J. Poplawski, Big bounce and closed universe from spin and torsion, *The Astrophysical Journal* **870**, 78 (2019), arXiv:1808.08327 [gr-qc].
- [59] J. I. Kapusta and C. Gale, *Finite-Temperature Field Theory: Principles and Applications*, 2nd ed., Cambridge Monographs on Mathematical Physics (Cambridge University Press, 2006).
- [60] CMB-S4 Collaboration, K. Abazajian, *et al.*, CMB-S4 Science Book, First Edition, arXiv e-prints (2016), arXiv:1610.02743 [astro-ph.CO].
- [61] M. Brüggen, Effects of a torsion field on Big Bang nucleosynthesis, *General Relativity and Gravitation* **31**, 1935 (1999), arXiv:astro-ph/9906403 [astro-ph].
- [62] D. Brout *et al.*, The Pantheon+ analysis: Cosmological constraints, *Astrophys. J.* **938**, 110 (2022), arXiv:2202.04077 [astro-ph.CO].
- [63] DES Collaboration, T. M. C. Abbott, *et al.*, The dark energy survey: Cosmology results with ~ 1500 new high-redshift type ia supernovae using the full 5-yr data set, *Astrophys. J. Lett.* **973**, L14 (2024), arXiv:2401.02929 [astro-ph.CO].
- [64] S. Alam *et al.*, Completed SDSS-IV extended baryon oscillation spectroscopic survey: Cosmological implications from two decades of spectroscopic surveys at the Apache Point Observatory, *Phys. Rev. D* **103**, 083533 (2021), arXiv:2007.08991 [astro-ph.CO].
- [65] A. J. Ross, L. Samushia, C. Howlett, W. J. Percival, A. Burden, and M. Manera, The clustering of the SDSS DR7 main galaxy sample – I. a 4 per cent distance measure at $z = 0.15$, *Mon. Not. R. Astron. Soc.* **449**, 835 (2015), arXiv:1409.3242 [astro-ph.CO].
- [66] F. Beutler, C. Blake, M. Colless, D. H. Jones, L. Staveley-Smith, L. Campbell, Q. Parker, W. Saunders, and F. Watson, The 6dF Galaxy Survey: baryon acoustic oscillations and the local Hubble constant, *Mon. Not. R. Astron. Soc.* **416**, 3017 (2011), arXiv:1106.3366 [astro-ph.CO].
- [67] DES Collaboration, T. M. C. Abbott, *et al.*, Dark Energy Survey Year 3 results: Cosmological constraints from galaxy clustering and weak lensing, *Physical Review D* **105**, 023520 (2022), arXiv:2105.13549 [astro-ph.CO].
- [68] J. Torrado and A. Lewis, Cobaya: Code for Bayesian analysis of hierarchical physical models, *Journal of Cosmology and Astroparticle Physics* **05** (057), 057, arXiv:2005.05290 [astro-ph.IM].
- [69] P. Diego-Palazuelos, J. R. Eskilt, Y. Minami, M. Tristram, *et al.*, Cosmic birefringence from the Planck data release 4, *Phys. Rev. Lett.* **128**, 091302 (2022), reports beta = 0.30 +/- 0.11 deg from Planck NPIPE (PR4), arXiv:2201.07682 [astro-ph.CO].
- [70] D. Alonso, J. Sanchez, and A. Slosar (LSST Dark Energy Science), A unified pseudo- C_ℓ framework, *Mon. Not. Roy.*

- Astron. Soc. **484**, 4127 (2019), arXiv:1809.09603 [astro-ph.CO].
- [71] T. Fujita, K. Murai, H. Nakatsuka, and S. Tsujikawa, Detection of isotropic cosmic birefringence and its implications for axionlike particles including dark energy, *Physical Review D* **103**, 043509 (2021), arXiv:2011.11894 [astro-ph.CO].
- [72] M. Walmsley, C. Lintott, T. Géron, *et al.*, Galaxy Zoo DECaLS: Detailed visual morphology measurements from volunteers and deep learning for 314 000 galaxies, *Monthly Notices of the Royal Astronomical Society* **509**, 3966 (2022), arXiv:2102.08414 [astro-ph.GA].

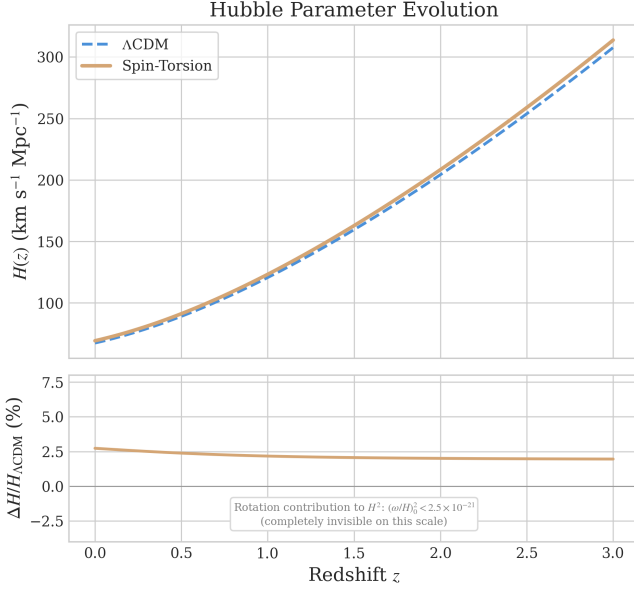


FIG. 3. ECH dark-energy model vs. Λ CDM Hubble evolution. Upper panel (y -axis: $H(z)$ [$\text{km s}^{-1} \text{Mpc}^{-1}$]): Hubble parameter for the full ECH dark-energy model (orange; ΞM_{Pl}^2 term from Eq. 10) versus Λ CDM (blue). Lower panel (y -axis: $\Delta H/H_{\Lambda\text{CDM}}$ [%]): percent deviation between the spin-torsion benchmark cosmology and a Planck-VI Λ CDM reference; an illustrative parameter-set comparison under the ΞM_{Pl}^2 phenomenological on-shell scaling ansatz (Appendix B), not a derived prediction. The Lite-BIRD/SPHEREx falsification windows (§III, Sec. X G; 2027–early-2030s) carry the load-bearing testable predictions. The rotation contribution $c_\omega \omega^2$ is a *distinct* and negligible term, confined to $\lesssim 10^{-21} \rho_\Lambda^{\text{obs}}$ ($(\omega/H)_0^2 < 2.5 \times 10^{-21}$; dividing by $3\Omega_\Lambda \approx 2.1$ gives $\sim 1.2 \times 10^{-21}$ of $\rho_\Lambda^{\text{obs}}$) — completely invisible on the scale plotted. The dark-energy mechanism is therefore the ΞM_{Pl}^2 term sourced by the parity-odd contorsion sector (§II A 2), *not* the rotation component. The orange ECH curve uses Ξ set to reproduce ρ_Λ (i.e. $\Xi = \rho_\Lambda/M_{\text{Pl}}^4 = \Lambda_{\text{eff}}/M_{\text{Pl}}^2 \approx 10^{-123}$, consistent with the dimensionless-ratio derivation in the body below) with spin-torsion benchmark cosmology $H_0 = 69.2 \text{ km/s/Mpc}$, $\Omega_m = 0.310$, and enhanced radiation density $\Omega_r^{\text{ext}} = \Omega_r^{\text{std}}(1 + 0.3\frac{7}{8}(\frac{4}{11})^{4/3})$ as a ΔN_{eff} proxy for the extra-radiation ECH channel; the Λ CDM reference uses $H_0 = 67.36 \text{ km/s/Mpc}$, $\Omega_m = 0.315$ (Planck-VI best-fit). The $\Delta H/H_{\Lambda\text{CDM}}$ deviation is $\sim 2\text{--}3\%$ across $z = 0\text{--}3$, but *this deviation is dominated by the differing H_0 baselines, not by spin-torsion dynamics*: because $H(0) \equiv H_0$ for both models, the $z = 0$ residual is exactly the H_0 offset, $(69.2 - 67.36)/67.36 \approx 2.7\%$, independent of the ΞM_{Pl}^2 sector. An H_0 -matched comparison (identical H_0, Ω_m for both curves) isolates the genuine dynamical effect of the spin-torsion dark-energy term, a sub-percent residual sourced by the Ω_m and ΔN_{eff} differences alone; the figure is thus an illustrative benchmark-parameter overlay, *not* a measure of the spin-torsion signal, and the 2–3% figure must not be read as a bounce signature. The benchmark $H_0 = 69.2$ used here is a deliberately high illustrative value and differs from the paper’s adopted $H_0 = 67.68 \pm 1.06$ (Table I, Table VI). The $(\omega/H)_0 < 5 \times 10^{-11}$ bound (Saadeh *et al.* [21]) applies under the Bianchi IX isotropized cosmological model and is adopted here as a conservative bookkeeping upper limit on $c_\omega \omega^2$.




FIG. 4. **Observational decision timeline for the two surviving ECH-independent class-level falsification paths.** Top: LiteBIRD CMB birefringence ($\sigma(\beta) \approx 0.03^\circ$, launch early 2030s) testing the spectator-ALP route 4. Bottom: SPHEREx galaxy bispectrum (~ 2028 first cosmological data release) testing the matter-bounce $f_{NL} = -35/16$ prediction at $1.3\text{--}2.75\sigma$ realistic significance (footnote 6). Both surveys deliver ECH-independent class-level discrimination tests; under the stated ansätze they can falsify the relevant matter-bounce and uniform spectator-ALP benchmarks, but they do not identify a unique surviving minimal-ECH channel. Curve families labeled $\rho = 0, 0.3, 0.5$ correspond to the assumed cross-correlation coefficient ρ between the f_{NL} and β joint-forecast estimators (with $\rho = 0$ the uncorrelated baseline; $\rho > 0$ combinations track the gain in joint significance under a positive between-estimator correlation).




FIG. 5. **Parameter-naturalness diagnostics for the minimal-ECH dark-energy parameterization.** *Top* (y -axis: α/M [GeV^{-1}]): renormalization-group running of the parity-odd coupling from the present epoch to the Planck scale, anchored at the primordial benchmark $\alpha/M \sim 10^{-21} \text{ GeV}^{-1}$ (Sec. IV F). *Bottom* (y -axis: fine-tuning score [orders of magnitude]): comparison of residual tuning (unreduced M_{Pl} convention throughout): Λ CDM (10^{122}), quintessence (10^{60}), $f(R)$ gravity (10^{40}), and the spin-torsion N_{tot} parameterization of this work (10^5). The Λ CDM (10^{122}) and spin-torsion (10^5) scores are derived in this paper (Appendix B and Sec. XII A); the quintessence (10^{60}) and $f(R)$ (10^{40}) entries are illustrative order-of-magnitude literature-level comparators, not derived here. The 10^5 residual annotation is the score under the N_{tot} reparameterization; per Sec. XII A this is a reparameterization of the cosmological-constant problem as sensitivity to N_{tot} , *not* a resolution. All four enumerated routes either sit outside their naturalness windows or require a $m_\theta \sim H_0$ tuning that re-imports the cosmological-constant problem; the 13 mechanism-class constraints (Table IV) constrain the routes at the amplitude level. (σ values across panels use different null procedures; see text.)

7 Foundations (A–G): mechanism-class barriers




FIG. 6. **Structure of the 14-barrier closure.** The 7 foundation mechanism classes (Foundations A–G, top row; Barriers 1–7) and 6 observational-channel branches (Branches H, J, L, M, N, O, middle row; Barriers 8–14) feed collectively into the four closed dark-energy routes R1–R4 (bottom row). Arrows indicate which barrier classes constrain each route. B8 (Branch H, parity-even interaction) is subsumed by B14 (perturbation transparency) and they are grouped together. Thirteen mechanism-class constraints are shown (several share the scaling ansatz but each probes a distinct physical failure mode); B8 is not counted separately.




FIG. 7. **Detection forecast for the two surviving ECH-independent class tests.** Top: matter-bounce $f_{\text{NL}} = -35/16$ in the SPHEREx multi-tracer f_{NL} Fisher landscape (companion [2], $1.3\text{--}2.75\sigma$ projection). Bottom: spectator-ALP cosmic birefringence in the LiteBIRD $\sigma(\beta) \approx 0.03^\circ$ window (Appendix E); the WMAP+Planck $\beta = 0.342^\circ \pm 0.094^\circ$ and ACT DR6 $\beta = 0.215^\circ \pm 0.074^\circ$ points are shown for reference. The SPHEREx f_{NL} forecast is potentially decisive in optimistic configurations; $1.3\text{--}2.75\sigma$ after the stated systematic budget (Table I footnote b) against $f_{\text{NL}} = 0$; the LiteBIRD birefringence forecast targets a non-zero- β detection at its $\sigma(\beta) \approx 0.03^\circ$ sensitivity but does not separate $\beta = 0.27^\circ$ from the published $\beta = 0.342^\circ \pm 0.094^\circ$ at high significance (the discrimination runs at $\sim 0.7\sigma$ given current central values); neither is uniquely an ECH prediction (§XIII). The significance tracks duplicate the $\rho = 0$ combination of Fig. 4 (ρ here is the cross-correlation coefficient between the f_{NL} and β joint-forecast estimators; see the Fig. 4 caption for the precise definition); this panel is the summary view retained for the surviving-tests discussion, while Fig. 4 additionally shows the correlated ($\rho = 0.3, 0.5$) combinations and milestone annotations. (σ values across panels use different null procedures; see text.)




FIG. 8. Full-tension MCMC corner plot (119,617 post-burnin samples, `getdist`-thinned from 176,240 raw; footnote 9) over Planck+BAO+SN+H0+S₈. The ΔN_{eff} posterior is consistent with zero (-0.020 ± 0.169), confirming no additional relativistic species at recombination.




FIG. 9. ΔN_{eff} marginal posterior comparison across the two frozen dataset combinations of Table VII (176,240 and 132,949 samples). Panel (a): Gaussian summaries of the ΔN_{eff} marginal posteriors at the Table VII means $\pm 1\sigma$, with the Standard-Model value $\Delta N_{\text{eff}} = 0$ marked. Panel (b): all seven Table VII parameters, normalized to the full-tension mean and σ . Both combinations recover ΔN_{eff} consistent with zero at $< 1\sigma$; the headline “full-tension” stack (Planck+BAO+SN+H₀+S₈) gives $\Delta N_{\text{eff}} = -0.020 \pm 0.169$. No evidence for a recombination-era ΔN_{eff} shift appears in this stock-CAMB proxy run; this does *not* directly test the ECH spin-torsion sector, which lacks a Boltzmann-module prediction for ΔN_{eff} in stock-CAMB (see § E 2 scope note).

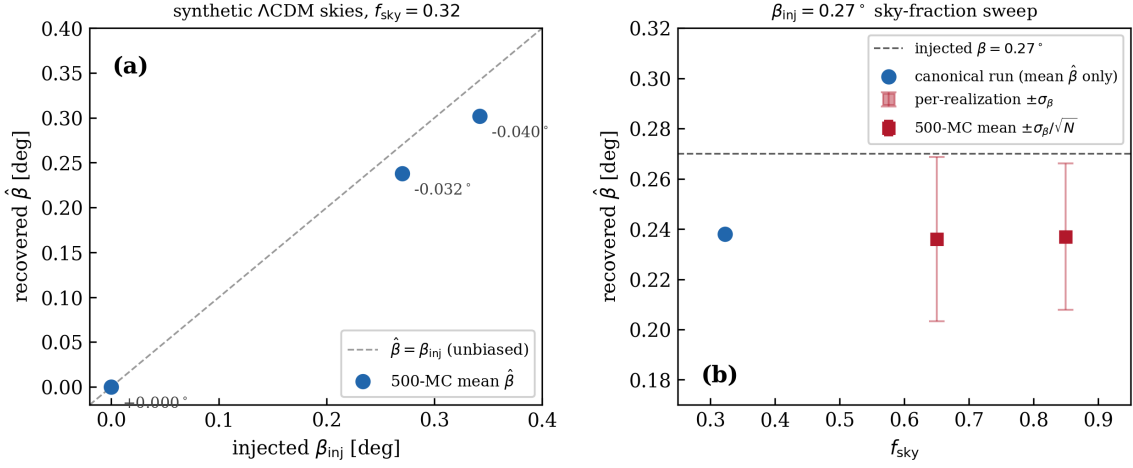


FIG. 10. NaMaster pipeline-recovery validation on synthetic Λ CDM polarization skies ($N_{\text{side}} = 512$, $\Delta_P = 10 \mu\text{K} \cdot \text{arcmin}$ white noise, $N = 500$ MC realizations per point; Sec. F 1). Panel (a): 500-MC mean recovered $\hat{\beta}$ vs. injected $\beta_{\text{inj}} \in \{0, 0.27^\circ, 0.342^\circ\}$ at the canonical apodized $f_{\text{sky}} = 0.32$ ACT-like mask, annotated with the pipeline-recovery bias $\hat{\beta} - \beta_{\text{inj}}$ (0.000° , -0.032° , -0.040°). Panel (b): the $\beta_{\text{inj}} = 0.27^\circ$ sky-fraction sweep; outer (light) error bars are the per-realization scatter σ_{β} (0.029° at $f_{\text{sky}} = 0.85$, 0.033° at 0.65), inner bars the standard error of the 500-MC mean σ_{β}/\sqrt{N} ; per-realization σ_{β} was not recorded in the original canonical $f_{\text{sky}} = 0.32$ artifact, so that point is plotted with the mean only; a dedicated 500-MC rerun (fn. 12) measures $\sigma_{\beta} = 0.046^\circ$ at this point. The worst-case $|\text{bias}| = 0.040^\circ \pm 0.002^\circ$ (SE of 500-MC mean; $\sigma_{\beta} = 0.046^\circ$ at $f_{\text{sky}} = 0.32$; `namaster_500mc.py`) is carried forward as the observed NaMaster pipeline bias (deconvolution-algebra bias on foreground-free skies; *estimator-specific* — this figure is tied to the unweighted χ^2 template fit chosen to match published scripts, and the inverse-variance-weighted estimator removes $\approx 80\%$ of it (Sec. F 1); it is *not* a real-sky systematics bound, which additionally requires foreground cleaning and breaking of the β - α degeneracy by unrotated galactic foregrounds). Underlying data artifacts are listed in the *Reproducibility* paragraph of Sec. F 1.




FIG. 11. Spectator-ALP joint posterior triangle from the continuous-prior cross-check configuration in which the photon anomaly coefficient is sampled freely (flat priors $C_{a\gamma} \in [4, 60]$ — shifted and extended from the earlier [1, 30] to cover the posterior-supported coupling band (median 20.7, 16–84% [7.3, 45.6]); the dropped [1, 4] interval lies entirely below the minimum coupling ≈ 8.6 required to reach the *central* value $\beta = 0.342^\circ$ at the largest in-box displacement $\Delta\phi/f_a = 1.19$ — couplings in [4, 8.6) remain posterior-supported because the summary likelihood permits β below the central value — $\theta_i \in [0.01, \pi]$, $\log_{10}(m_a/\text{eV}) \in [-35, -30]$ — the mass prior corresponds to $m/H_0 \approx 7 \times 10^{-3}$ to 7×10^2 for $H_0 = 67.7 \text{ km/s/Mpc} = 1.44 \times 10^{-33} \text{ eV}$), broadening the fixed- $C_{a\gamma} = 8$ and $C_{a\gamma} \in [1, 30]$ configurations of Appendix H 3; the direct-sample priors are on the underlying ALP parameters (θ_i, m_a), with $\Delta\phi/f_a$ derived along each trajectory. The rotation marginal, $\beta = 0.326^\circ \pm 0.099^\circ$, is consistent within 1σ with the fixed- $C_{a\gamma} = 8$ result $\beta_{\text{ALP}} = 0.336^\circ \pm 0.10^\circ$ and the observed $\beta_{\text{obs}} = 0.342^\circ \pm 0.094^\circ$ (an internal-consistency statement: all three are constrained by the same single published β_{obs} , not independent measurements). The θ_i – $C_{a\gamma}$ anti-correlation band traces the constant-product degeneracy enforced by the summary-likelihood constraint through $C_{a\gamma} (\Delta\phi/f_a) \approx 10.3$, and the m_a marginal piles toward the upper (heavier) edge of its prior range — the data-driven pull toward the upper edge of the natural-prior box discussed in the text, consistent with the model *accommodating* rather than naturally predicting the observed rotation (fn. 14).