

Technical Reproducibility and Consistency-Check Companion to the ECH Spin-Torsion Program:

Λ CDM+ ΔN_{eff} MCMC Proxy, NaMaster Pipeline Recovery, and a Birefringence Consistency Check with a Spectator-ALP Model

Houston Golden^{1,*}

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

(Dated: July 6, 2026)

We report the technical reproducibility and consistency-check material for the Einstein-Cartan-Holst (ECH) spin-torsion cosmology no-go program of Paper I(a) [1]. *Scope, stated up front:* none of the three analyses below implements or tests a torsion-modified Boltzmann/theory module, and none *verifies* the ECH spin-torsion sector. Each is an adjacent numerical cross-check — a stock-CAMB generic-radiation null proxy, a foreground-free synthetic-pipeline recovery validation, and a literature-data ALP accommodation — so the paper’s contribution is precisely a reproducibility and null-consistency note for three limited numerical cross-checks, not independent evidence for (or verification of) the spin-torsion theory. What this companion nonetheless establishes on its own is concrete and self-contained: (i) a *first-principles* derivation (Sec. III A) that the minimal ECH spin-torsion sector’s radiation-era contribution is Planck-suppressed to $\Delta N_{\text{eff}}^{(\text{ECH})} \sim 10^{-44}$ at BBN — some forty orders of magnitude below any foreseeable sensitivity, and a genuine result of *this* paper rather than one imported from Paper I(a); (ii) a validated NaMaster $E \rightarrow B$ pseudo- C_ℓ recovery pipeline with a quantified synthetic-sky bias budget; (iii) a quantified prior-predictive accommodation cost for the spectator-ALP birefringence fit; and (iv) a complete public reproducibility manifest backing all three analyses. Three analyses are documented. (1) *Stock-CAMB Λ CDM+ ΔN_{eff} MCMC proxy* (Cobaya v3.6.1, **309,189** frozen samples across two converged dataset combinations): this run uses stock CAMB with ΔN_{eff} as a free parameter (consistent with the scope stated above — a null-consistency test of an extra radiation-like degree of freedom). Both frozen dataset combinations find ΔN_{eff} consistent with zero (-0.020 ± 0.169 full-tension; $+0.058 \pm 0.179$ Planck+BAO+SN) and H_0 consistent with the standard Planck- Λ CDM value (67.68 ± 1.06 full-tension; 67.78 ± 1.09 Planck+BAO+SN, both in $\text{km s}^{-1} \text{Mpc}^{-1}$; the ΔN_{eff} extension does *not* reduce the residual $\sim 3.6\sigma$ tension with the SH0ES local-distance-ladder $H_0 = 73.04 \pm 1.04 \text{ km s}^{-1} \text{Mpc}^{-1}$). We additionally derive, from first principles, the bespoke ECH-sector contribution to ΔN_{eff} : integrating out the algebraic torsion yields a Planck-suppressed four-fermion interaction whose radiation-era energy density scales as $(T/M_{\text{Pl}})^2$, giving $\Delta N_{\text{eff}}^{(\text{ECH})} \sim 10^{-44}$ at BBN — a strictly negligible value. The stock-CAMB run is thus not a proxy for an unknown prediction but a ~ 40 -orders-of-magnitude-conservative observational envelope on this bespoke ECH result (Sec. III A). (2) *NaMaster pseudo- C_ℓ pipeline validation* on synthetic Λ CDM CMB polarization skies with an ACT-like footprint mask ($N_{\text{side}} = 512$, $\ell_{\text{max}} = 1024$, $f_{\text{sky}} = 0.32$, $10 \mu\text{K} \cdot \text{arcmin}$ white noise, 500 Monte Carlo realizations): injecting the spectator-ALP fiducial value $\beta = 0.27^\circ$ recovers $\hat{\beta} = 0.238^\circ$ (pipeline-recovery bias $\hat{\beta} - \beta_{\text{inj}} = -0.032^\circ$; the worst-case bias across injections, -0.040° at $\beta_{\text{inj}} = 0.342^\circ$, is carried forward as the foreground-free synthetic-pipeline recovery bias — both are MC pipeline-recovery figures, not sky-measurement systematics, and are not directly comparable to each other’s published sky significances). *Scope of the validation:* the test confirms the algebraic pseudo- C_ℓ $E \rightarrow B$ deconvolution under MASTER [2] mode coupling, NOT the physical separation of the cosmic-rotation angle β from the instrumental-miscalibration angle α which strictly requires unrotated galactic foregrounds (the synthetic CMB-only skies contain no galactic foregrounds, so the very component that breaks the β - α degeneracy in published Planck/ACT DR6 measurements is absent by construction). The MC recovery is therefore a pipeline-validation figure, not a sky-detection significance claim. The primary sky detection significance is the published Planck/ACT DR6 2.7 – 2.9σ [3, 4] (the $\beta = 0.342^\circ \pm 0.094^\circ$, 3.6σ headline used throughout this paper is from the published PR3+WMAP9 joint analysis of Eskilt & Komatsu [5]; the PR3-vs-PR4/NPIPE disambiguation is given in fn. 3 in Sec. IV); the pipeline SNR figures refer to recovery of injected MC signals and are *not* competitive sky measurements. (3) *Spectator-ALP consistency check:* for a field with $f_a \sim M_{\text{Pl}}$ (where $M_{\text{Pl}} = (8\pi G)^{-1/2} \approx 2.44 \times 10^{18} \text{ GeV}$ is the reduced Planck mass), the scan-prior $m \sim H_0$ region brackets the published joint WMAP+Planck signal $\beta = 0.342^\circ \pm 0.094^\circ$ (3.6σ) [5], but the posterior-supported fixed-coupling ($C_{a\gamma} = 8$) accommodation shifts to $m \gg H_0$ (median $m \simeq 36 H_0$ at fixed $C_{a\gamma} = 8$), and the spectator-safe ($\Omega_a < 0.01$) subset is tuned (see Table IV for the per-subset readout). *Prior-predictive accommodation cost:* to test directly whether this consistency is tautological — i.e. whether the model reproduces β_{obs} merely because the coupling and misalignment priors are tuned to it — we run a prior-predictive Monte-Carlo, drawing the ALP nuisance parameters from their stated priors, propagating each draw through the committed equation of motion with *no* likelihood weighting, and counting the fraction that lands within the Eskilt et al.

band ($\beta_{\text{obs}} = 0.342^\circ \pm 0.094^\circ$; artifact [reproducibility/cosmology/alp_prior_predictive.py](#), $N = 100,000$ draws per configuration). At the posterior-supported fixed coupling $C_{a\gamma} = 8$, 11.6% of prior draws fall within 1σ of β_{obs} (23.9% within 2σ); under the broader continuous-coupling prior $C_{a\gamma} \sim \mathcal{U}[4, 60]$ these drop to 6.1% and 12.6%. This *moderate* fraction cuts against the tautology framing in both directions: a genuinely tautological accommodation — one that could reproduce *any* target β — would give an order-unity fraction, whereas a knife-edge fine-tuning would give $\ll 1\%$; the observed one-in-eight-to-one-in-twenty draws instead places the birefringence signal in a reachable but non-generic region of the ALP prior. We therefore report the check as a genuine (if not parameter-free) accommodation rather than a tuned fit, and disclose the fraction explicitly. *Spectator-status caveat*: for $\theta_i \sim 1$ the ALP energy density $\rho_a \sim m^2 f_a^2 \theta_i^2 \sim H_0^2 M_{\text{Pl}}^2$ is of order the critical density today, so the spectator label is only consistent under the explicit restriction $\theta_i \ll 1$ (which constitutes a $\gtrsim 100$ -fold fine-tuning of the misalignment initial condition under the natural $\cos\theta_i$ -flat prior — equivalently ~ 25 -fold relative to the ad-hoc $\theta_i \approx 0.5$ midpoint — quantified in Sec. VI); at $\theta_i \sim 1$ the ALP must instead be treated as the dark-energy field itself, which is outside the scope of this companion-paper consistency check. The same birefringence arises in standard GR with an identical ALP; it is not a distinctive ECH prediction, and the spectator-status fine-tuning is required regardless of whether the underlying cosmology is a bounce or Λ CDM. All three analyses are compatibility checks only; no analysis here constitutes evidence for or against a bounce cosmology, and all model-preference inference (Bayes factor, $\Delta\text{AIC/BIC}$, nested-sampling evidence) is deferred to follow-up nested-sampling work. A reproducibility manifest is included in Appendix A.

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

CONTENTS

I. Introduction	2
II. Cosmological Tensions: H_0 and σ_8	4
III. Stock-CAMB Λ CDM+ ΔN_{eff} MCMC: Generic Radiation-Proxy Test (Not a Spin-Torsion Theory Module)	4
A. Bespoke ECH-sector ΔN_{eff} from the torsion-induced four-fermion interaction	4
IV. Data Methods: CMB E - B Analysis	7
V. Cosmological Fits and Model Comparison	11
A. Datasets and Configuration	11
B. Results: Λ CDM+ ΔN_{eff} proxy	12
VI. Cosmic Birefringence: Spectator ALP Consistency Check	12
ALP dark-energy fraction Ω_a : definition and computation	16
VII. Conclusions	16
Data and Code Availability	17
Acknowledgments	19
A. Reproducibility Materials	19
B. Claims Classification	20
C. ALP-MCMC Sampled Parameters, Priors, and Likelihood Stack	20

References

21

I. INTRODUCTION

This companion paper provides the technical reproducibility and consistency-check layer for the ECH structural-closure no-go result reported in Paper I(a) [1]. We use “companion” in the strict sense: the paper documents three limited numerical cross-checks and their reproducibility artifacts, and it does *not* implement, test, or verify any torsion-modified component of the ECH theory itself. Its standing relative to Paper I(a) is that of a null-consistency and reproducibility note, not of independent confirmation. That said, the present paper does contribute one genuinely new first-principles result of its own — the derivation (Sec. III A) that the minimal ECH spin-torsion sector’s radiation-era ΔN_{eff} is Planck-suppressed to $\sim 10^{-44}$ at BBN, some forty orders of magnitude below any foreseeable sensitivity — alongside a validated $E \rightarrow B$ recovery pipeline, a quantified spectator-ALP accommodation cost, and a complete public reproducibility manifest, all self-contained within this manuscript. Paper I(a) establishes channel-level closure of four independent minimal-ECH dark-energy routes (kinetic-coupling channel, potential-energy channel, perturbation-transparency channel, and spin-torsion propagator channel) and proves a perturbation-transparency theorem for the Holst sector; the combined structural analysis closes the minimal-ECH parameter space as a dark-energy scenario and motivates the numerical cross-checks of this companion. The present paper documents the three numerical analyses that support and contextualize those results. This is a coordinated concurrent submission: Paper I(a) supplies the physical motivation and the structural no-go result, while the three

* houston@hubify.com

TABLE I. Principal symbols used in the birefringence / spectator-ALP consistency check (Sec. VI) and the Λ CDM+ ΔN_{eff} proxy run (Secs. II, III). All quantities are defined at their first use in the text; this table collects them for reference. No quantity here is an ECH-derived prediction — see the up-front scope statements.

Symbol	Meaning
$C_{a\gamma}$	Dimensionless axion–photon coupling constant (KSVZ/DFSZ benchmark $\mathcal{O}(1)$; posterior-supported value 8)
β	Cosmic-birefringence rotation angle (deg); β_{obs} published datum, β_{ALP} the ALP-model posterior value
θ_i	ALP misalignment initial condition (dimensionless); natural $\cos \theta_i$ -flat prior, midpoint $\theta_i \approx 0.5$
Ω_a	Present-day ALP energy-density fraction of critical density, $\rho_a(z=0)/\rho_{\text{crit},0}$; spectator cut $\Omega_a < 0.01$
f_a	ALP decay constant (GeV); fiducial $f_a \sim M_{\text{Pl}}$
m	ALP mass, quoted in units of H_0 ; scan-prior region $m \sim H_0$
M_{Pl}	Reduced Planck mass, $(8\pi G)^{-1/2} \approx 2.44 \times 10^{18}$ GeV
ΔN_{eff}	Extra effective relativistic d.o.f. (generic radiation proxy in the MCMC; ECH bound in Sec. III A)
α	Instrumental polarization-miscalibration angle (deg), degenerate with β absent unrotated galactic foregrounds

numerical analyses below — and their data products, chains, and pipeline code — are self-contained and independently reproducible from this paper’s own repository, so no quantitative claim here depends on unpublished companion numerics. Because the two papers are posted together, the companion Paper I(a) arXiv identifier is inserted in place of the marker [arXiv:XXXX.XXXXX] wherever Paper I(a) is cited at coordinated submission (and, reciprocally, this paper’s identifier is inserted into Paper I(a)); until then the marker stands as a deliberate placeholder.

Imported theory results (summarized for standalone readability).—So that the present companion can be read without Paper I(a) in hand, we restate the load-bearing structural results that Paper I(a) establishes and that this paper takes as fixed theoretical inputs; the full derivations are in Paper I(a), posted concurrently, and are not reproduced here. (i) *Thirteen mechanism-class structural barriers.* Paper I(a) catalogs thirteen distinct mechanism classes (organized in its barrier table) by which the minimal Einstein–Cartan–Holst (ECH) spin-torsion sector fails to source a viable late-time dark-energy component; each barrier closes a separate route, and collectively they close the minimal-ECH parameter space as a dark-energy scenario. (ii) *Perturbation-transparency theorem.* For the Holst sector, Paper I(a) proves that the torsion/Holst contributions are transparent to linear cosmological perturbations—they do not propagate as independent perturbative degrees of freedom—so the “perturbation-transparency” channel (one of the four minimal-ECH dark-energy routes: kinetic-coupling, potential-energy, perturbation-transparency, and spin-torsion-propagator) closes already at the background level and leaves no distinctive linear-perturbation dark-energy signature. (iii) *Surviving matter-bounce signature* $f_{\text{NL}} = -35/8$. With the dark-energy channels closed, the distinctive observational target that survives in Paper I(a) is the matter-bounce non-Gaussianity amplitude $f_{\text{NL}} = -35/8$, established structurally there and valid strictly for the minimal scalar-only $w = 0$ matter-dominated contraction class (not the broader bouncing-cosmology landscape—ekpyrotic,

Cuscuton, string-gas, and quintom variants carry different f_{NL} signatures). This signature is quoted here only as program context: it is *not* probed by the present companion, whose two analyses test ΔN_{eff} and birefringence consistency. The companion’s own contribution—the MCMC posteriors, pipeline-validation figures, and spectator-ALP consistency check reported below—is self-contained and independently reproducible from this paper’s repository.

Scope of this paper.—Three numerical cross-check analyses are reported here, each a compatibility/consistency check rather than a verification of the ECH sector, and each with explicit scope limitations:

1. Λ CDM+ ΔN_{eff} MCMC proxy (Secs. II and III): stock CAMB Boltzmann solver with ΔN_{eff} as a free parameter. **Not a spin-torsion theory module** (see the §III scope statement); this run tests whether the data prefer an extra radiation-like degree of freedom in standard cosmology.
2. *NaMaster CMB E-B analysis* (Sec. IV): a bias-injection Monte Carlo validation of the pseudo- C_ℓ deconvolution pipeline on synthetic Λ CDM polarization skies. **Not a competitive sky detection.** The high pipeline template-fit SNR figures (e.g., 20.32) refer to recovery of injected MC signals, not to the significance of the CMB sky measurement, which remains the published 2.7–2.9 σ from Planck/ACT.
3. *Spectator-ALP consistency check* (Sec. VI): a standard GR+ALP computation showing that the observed signal is consistent with an ALP 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 $\simeq 36 H_0$), well above the prior midpoint. A spectator-safe interpretation requires a tuned misalignment subspace (the $\sim 25\times$ misalignment tuning required to reconcile the observed signal with the spectator-consistent corner is disclosed in Sec. VI and fn. 6). **Not a distinctive ECH prediction.** The same $\beta \approx 0.27^\circ$ arises in any GR+ALP setup with the same parameters; no ECH-specific derivation connects the Holst action to the photon-torsion coupling

required.

What is NOT in this paper.—The full derivations of the 13 mechanism-class structural barriers, the perturbation-transparency theorem, the 14-barrier table, and the surviving matter-bounce-specific test predictions ($f_{\text{NL}} = -35/8$, valid strictly for the minimal scalar-only $w = 0$ matter-dominated contraction class, not the broader bouncing-cosmology landscape which encompasses ekpyrotic, Cuscuton, string-gas, and quintom variants with different f_{NL} signatures, and ALP birefringence) are in Paper I(a) [1] (those results are summarized for standalone context above). The SPHEREx multi-tracer Fisher forecast (in preparation, [6]; results to be detailed separately) is the subject of Paper II. The multi-survey anomaly catalog (in preparation, [7]) is the subject of Paper III. The galaxy chirality catalog (in preparation, [8]) is the subject of Paper IV.

Cross-paper citations.—When this companion reports MCMC values (H_0 , σ_8 , etc.) that are referenced in the main paper, those values come from Secs. III and V here.

II. COSMOLOGICAL TENSIONS: H_0 AND σ_8

The bounce scenario motivates extending Λ CDM 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. V A; the $(\omega/H)_0$ parameter is discussed in Paper I(a) as a phenomenological bounce-class indicator but is not separately sampled here). The full-tension dataset combination includes the SH0ES [9] 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.

III. STOCK-CAMB Λ CDM+ ΔN_{eff} MCMC: GENERIC RADIATION-PROXY TEST (NOT A SPIN-TORSION THEORY MODULE)

Scope statement.—“MCMC verification” refers throughout to a stock CAMB run of Λ CDM 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. III 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 [10]. Integrating torsion out of the action leaves an effective four-fermion contact interaction of Nambu–Jona-Lasinio type [11, 12],

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

where γ is the Barbero–Immirzi parameter of the Holst term [12, 13]; in the pure Einstein–Cartan–Sciama–Kibble limit $\gamma \rightarrow \infty$ the bracket $\rightarrow 1$ and Eq. (1) reduces to the canonical $-\frac{3}{2}\pi G_N (\bar{\psi}\gamma^5\gamma^\mu\psi)^2$ [11]. 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. (1) 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 [14]. Explicitly, the energy density contributed by the dimension-6 contact operator of Eq. (1) 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 [15], with $\mathcal{O}(1)$ coefficients dropped

rather than a rigorous NJL finite- T computation), so that

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

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 [14]. Its contribution to the effective radiation budget at the epochs probed by the data is

$$\Delta N_{\text{eff}}^{(\text{ECH})} \sim \left(\frac{T}{M_{\text{Pl}}} \right)^2 \sim \begin{cases} 1.7 \times 10^{-43}, & T \sim 1 \text{ MeV}, \\ 1.1 \times 10^{-56}, & T \sim 0.26 \text{ eV}, \end{cases} \quad (3)$$

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. 2), 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-S4 targets $\sigma(N_{\text{eff}}) \sim 0.03$ [16]]. 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 [17]. We therefore do *not* claim a specific non-zero torsion ΔN_{eff} ; the honest, first-principles outcome is the upper bound of Eq. (3).

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 $\Delta N_{\text{eff}} \lesssim 0.31$ –0.40 (Sec. V) exceeds the theoretical ECH value Eq. (3) 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. (2) 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.lowIE`

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 0.04 σ agreement in ΔN_{eff} across this likelihood substitution provides an empirical bound on pairing-induced bias at the quoted precision; see Sec. V *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 `nmu` 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.¹

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 [10, 13] is dimension-6 and M_{Pl}^{-2} -suppressed²: 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 [18] (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 consis-

tent 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 SH0ES [9] 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 $\sigma_{M_B} = 0.049$ 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 (Mb 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 ΛCDM 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 ($\sigma(N_{\text{eff}}) \sim 0.03$ [16]) will sharpen this constraint substantially relative to the current Planck+BAO+SN level.

Independent cross-validation.—Liu *et al.* [19] con-

¹ 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, $176,240 \times 0.7 + 132,949 \times 0.7 \approx \mathbf{216,432}$ post-burnin samples remain across both frozen chains (`convergence_summary.json`). For the full-tension subset specifically, $176,240 \times 0.7 \approx \mathbf{123,368}$ post-burnin (the **119,617** figure in Fig. 1 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 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 II (which contains only the two frozen combinations), and is *not* aggregated into the 309,189-sample headline anywhere in this paper.

² 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. [13], 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 II. 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. III 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 A). $S_8 \equiv \sigma_8 (\Omega_m/0.3)^{1/2}$, 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 ± 1.06	67.78 ± 1.09
ΔN_{eff}	-0.020 ± 0.169	$+0.058 \pm 0.179$
σ_8	0.803 ± 0.008	0.812 ± 0.009
S_8	0.814 ± 0.008	0.827 ± 0.010
Ω_m	0.308 ± 0.005	0.312 ± 0.006
τ	0.054 ± 0.007	0.056 ± 0.007
n_s	0.965 ± 0.006	0.967 ± 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 `HO_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 A).

strained an EC torsion model using DESI DR2 [20] + Pantheon+ [21] + DES-SN5YR [22] + 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. V, *Model-comparison statistics* paragraph), so no model-preference statement is made here that could agree or conflict with theirs.

IV. DATA METHODS: CMB E - B ANALYSIS

Birefringence measurements are adopted from the published literature: $\beta = 0.30^\circ \pm 0.11^\circ$ (Planck NPIPE [3]) and $\beta = 0.215^\circ \pm 0.074^\circ$ (ACT DR6 [4]). The spectator ALP analysis (Sec. VI) instead uses the Eskilt-Komatsu joint WMAP+Planck summary likelihood ($\beta = 0.342^\circ \pm 0.094^\circ$ [5])³ as its primary constraint; the Planck NPIPE and ACT DR6 values above are quoted only

³ Eskilt & Komatsu 2022 disambiguation: the published PRD paper [5] (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 likeli-

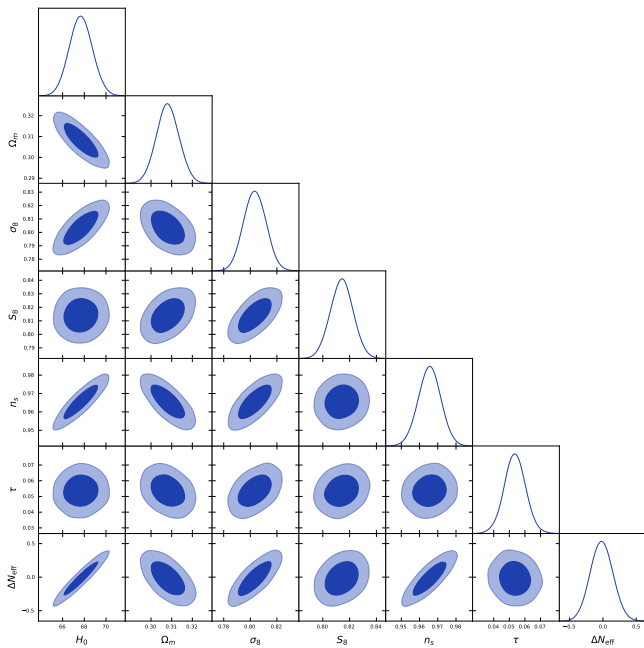


FIG. 1. Full-tension MCMC corner plot (119,617 post-burnin samples, `getdist`-thinned from 176,240 raw; footnote 1) over Planck+BAO+SN+H0+S₈. The ΔN_{eff} posterior is consistent with zero (-0.020 ± 0.169), confirming no additional relativistic species at recombination.

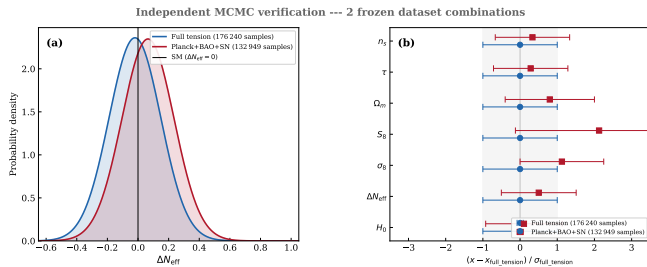


FIG. 2. ΔN_{eff} marginal posterior comparison across the two frozen dataset combinations of Table II (176,240 and 132,949 samples). Panel (a): Gaussian summaries of the ΔN_{eff} marginal posteriors at the Table II means $\pm 1\sigma$, with the Standard-Model value $\Delta N_{\text{eff}} = 0$ marked. Panel (b): all seven Table II 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 § III scope note).

hood $\beta = 0.342^\circ \pm 0.094^\circ$ 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.

for context and for the auxiliary inverse-variance cross-check of Eq. (8). (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 Eskiłt–Komatsu posterior shape.)

Scope note.—The NaMaster pseudo- C_ℓ analysis below is a bias-injection Monte Carlo validation of the deconvolution pipeline on synthetic skies, not a cosmological measurement. The high pipeline template-fit SNR figures (e.g., 20.32, 25.71; fn. 4) 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. Moreover, because the synthetic skies are CMB-only (no galactic foregrounds; see “Simulated skies” below), this validation cannot address the physical separation of the cosmic-rotation angle β from the instrumental-miscalibration angle α : the β – α degeneracy that dominates real-sky birefringence measurements is broken only by unrotated foregrounds, which are absent here by construction (cf. abstract).

Pipeline configuration.—The pseudo- C_ℓ analysis follows the NaMaster framework [23]. *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

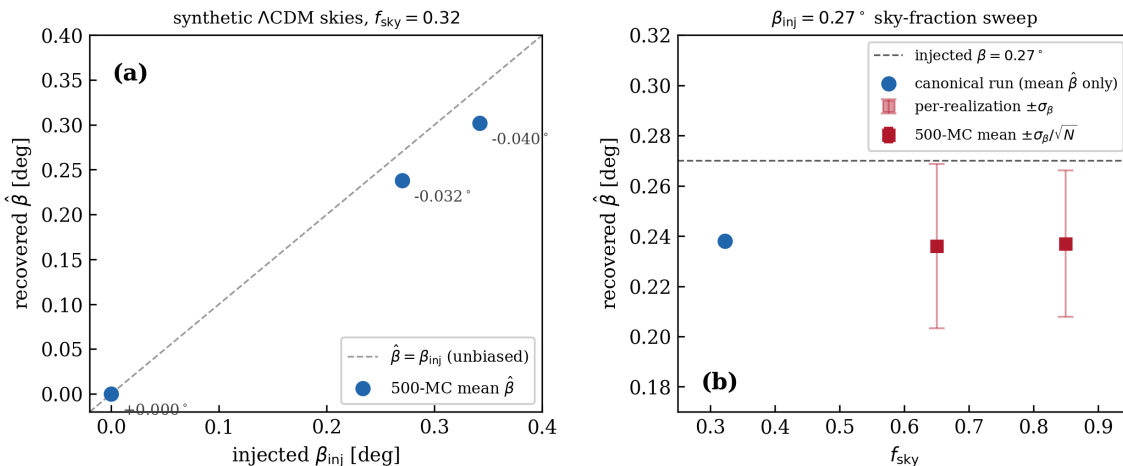


FIG. 3. NaMaster pipeline-recovery validation on synthetic Λ CDM polarization skies ($N_{\text{side}} = 512$, $\Delta_P = 10 \mu\text{K} \cdot \text{arcmin}$ white noise, $N = 500$ MC realizations per point; Sec. IV). Panel (a): 500-MC mean recovered $\hat{\beta}$ vs. injected $\beta_{\text{inj}} \in \{0, 0.27^\circ, 0.342^\circ\}$ 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. 4) 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. IV); 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. IV.

`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} 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 uni-

form 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

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

where $C_b^{EE, \text{tpl}}$ 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((cl_eb_measured - cl_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., [5]), 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 $C_b^{EE, \text{tpl}}$ 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 ($C_b^{EE, \text{templ}} = 0$ 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:

$$\hat{\beta}_{\text{NaMaster}} = 0.238^\circ \quad (500\text{-MC sample mean of } \hat{\beta}). \quad (5)$$

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 $\beta_{\text{obs}} = 0.342^\circ \pm 0.094^\circ$ [5]: 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 $C_\ell^{BB} = 0.05 C_\ell^{EE}$ 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).⁴ For $\beta = 0.342^\circ$ (the published joint WMAP+Planck value [5]), 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

⁴ 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, $\text{SNR}^{\text{templ}} \equiv [\sum_b (C_b^{EB, \text{th}} / \sigma_b^{\text{MC}})^2]^{1/2}$, where $C_b^{EB, \text{th}} = \sin(2\beta) \cos(2\beta) C_b^{EE}$ is the injected EB band-power 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 canonical $f_{\text{sky}} = 0.32$ 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 $\sqrt{N} = 22.4$.

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 sharpens 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 [5]; note $0.040^\circ \approx 0.5\sigma$ of the ACT DR6 $\sigma_\beta = 0.074^\circ$ [4], 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., [5]), 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/c10_robustness_battery.json`; `reproducibility/p1_namaster_500mc/results/c10_robustness_battery.json`.

V. COSMOLOGICAL FITS AND MODEL COMPARISON

A. Datasets and Configuration

We analyze four dataset combinations: (1) Planck CMB (NPIPE (PR4) CamSpec high- ℓ TTTEE + 2018 low- ℓ TT/EE + 2018 lensing) [24] (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) [25] + SDSS DR7 MGS [26] + 6dFGS [27]; (3) +Pantheon+ [21]; (4) +SH0ES M_B -anchor [9] + DES Y3 S_8 [28]. No DESI BAO likelihood enters the frozen Λ CDM+ ΔN_{eff} chains. The exact Cobaya likelihood blocks of all four chains are listed in Ta-

ble III. For scope clarity: Table II 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 III feed neither table. Parameter estimation uses Cobaya [29] (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. II). 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_lowlEE` and `planckpr4lensing`, differ from the frozen-chain names in Table III, `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/w0wa_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 re-run instead uses `planck_2020_lollipop_lowlEE` 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. III, limitation note).

VI. COSMIC BIREFRINGENCE: SPECTATOR ALP CONSISTENCY CHECK

Framing (prior-volume accommodation, not an independent test).—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 $\beta_{\text{obs}} = 0.342^\circ \pm 0.094^\circ$ [5], 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 $\Omega_a < 0.01$, requiring 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. 6) 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 IV) 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 §VI), 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 IV) 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

TABLE III. 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_highl_CamSpec.TTTEEE + planck_2018_lowl.TT + planck_2018_lowl.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 II)
Full-tension	+ <code>H0.riess2020Mb</code> + DES-Y3 S_8 Gaussian	frozen (Table II)

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. 6. 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.* [30].

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) [5] (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 [4], 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 background⁵ yields the field displacement from recombination to today, at the representative corner $m = 2H_0$, $\theta_i = 1$ of the natural-parameter box:

$$\Delta\phi/f_a \approx 0.42 \quad (m = 2H_0, \theta_i = 1). \quad (6)$$

⁵ 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 [31]. 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.

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`).⁶

Birefringence value.—For $C_{a\gamma} = 8$, $\theta_i = 1$, $m \approx 3.9 H_0$ (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):

$$\begin{aligned} \beta &\approx \frac{\alpha_{\text{EM}}}{4\pi} C_{a\gamma} \frac{\Delta\phi}{f_a} \\ &= \underbrace{(7.297 \times 10^{-3}/(4\pi))}_{\alpha_{\text{EM}}/(4\pi) \approx 5.81 \times 10^{-4}} \times 8 \times 1.06 \\ &= 4.93 \times 10^{-3} \text{ rad} \times \frac{180^\circ}{\pi} \approx 0.28^\circ. \end{aligned} \quad (7)$$

The product of the three factors is $5.81 \times 10^{-4} \times 8 \times 1.06 = 4.93 \times 10^{-3}$ 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

⁶ 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.

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 C 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. 6) 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 [3]) and $\beta = 0.215^\circ \pm 0.074^\circ$ (ACT DR6 [4]) via inverse-variance weighting:

$$\beta_{\text{combined}} = 0.241^\circ \pm 0.061^\circ \quad (\text{naive } 3.9\sigma \text{ upper bound}) \quad (8)$$

Note: this 3.9 σ is a deliberately-optimistic upper bound assuming zero Planck–ACT correlation; see Eskilt & Komatsu 3.6 σ for the properly-correlated headline. (Auxiliary cross-check only; the 3.9 σ 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 therefore *overestimate* the significance: the naive 3.9 σ figure is an upper bound on the true significance, not a lower bound. The published joint analysis at 3.6 σ [5], 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 3.9 σ 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 C) 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. 3; 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. 6 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 $\beta_{\text{obs}} = 0.342^\circ \pm 0.094^\circ$ [5] — 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 C; 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. 6); 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. 6) 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 C), 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 $\cos\theta_i$, 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 read-out* (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 IV), 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 IV, 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. 6 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. 4. 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 IV 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

$$V(\phi) = m_a^2 f_a^2 \left[1 - \cos\left(\frac{\phi}{f_a}\right) \right], \quad (9)$$

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

$$3H(z_{\text{osc}}) = m_a, \quad (10)$$

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 (10) 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:

$$\rho_a(z) = \rho_a(z_{\text{osc}}) \left(\frac{1+z}{1+z_{\text{osc}}} \right)^3, \quad z \ll z_{\text{osc}}. \quad (11)$$

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

$$\Omega_a \equiv \frac{\rho_a(z=0)}{\rho_{\text{crit},0}} \approx \frac{m_a^2 f_a^2 [1 - \cos(\theta_i)]}{\rho_{\text{crit},0} (1+z_{\text{osc}})^3}, \quad (12)$$

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. 6.

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. (10) for z_{osc} , and evaluates Eq. (12) 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. (12). 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 IV 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 $\sigma(\beta) \approx 0.03^\circ$ [32] (a forecast under foreground- and calibration-control assumptions, not a guaranteed instrument performance). For $\beta = 0.27^\circ$ this corresponds to $\sim 9\sigma$ statistical significance against the $\beta = 0$ null. This 9σ figure is *not* a model-discrimination forecast: separating the spectator-ALP value 0.27° from the current WMAP+Planck central value $0.342^\circ \pm 0.094^\circ$ is limited by the published measurement's uncertainty, $|0.342 - 0.27|/\sqrt{0.03^2 + 0.094^2} \approx 0.7\sigma$, so LiteBIRD alone will not separate the spectator-ALP value from the observed central value (the two figures correspond to distinct null hypotheses).

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. VI and fn. 6) and a non-minimal photon coupling; it does not uniquely predict the signal.

VII. CONCLUSIONS

This companion paper documents three technical reproducibility and consistency-check analyses for the ECH spin-torsion structural closure program [1]. To restate the scope plainly before the per-analysis summaries: none of the three *verifies* the ECH spin-torsion sector itself. The ΔN_{eff} run is a stock-CAMB generic-radiation null proxy (no torsion-modified Boltzmann equations); the NaMaster exercise is a foreground-free synthetic-pipeline recovery validation (not a sky measurement); and the ALP section is a literature-data accommodation of an already-published birefringence value (not an ECH-specific prediction). The paper's load-bearing output is therefore the set of frozen MCMC posteriors, the synthetic-pipeline recovery characteriza-

TABLE IV. 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 Eskill–Komatsu β_{obs} datum (Appendix C) 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)

tion, and the reproducibility artifacts; the ALP result is compatibility-only. All statements below should be read at that epistemic level — “consistent with”/“compatible with”/“does not exclude,” never “establishes” or “confirms the theory.”

Λ CDM+ ΔN_{eff} MCMC proxy.—The stock-CAMB Cobaya v3.6.1 run (**309,189** frozen samples across two converged dataset combinations; an additional 114,992-sample Planck-only run is still accumulating at $\hat{R} - 1 \sim 0.05$, is *not* reported in Table II, and is *not* aggregated into the frozen headline) finds ΔN_{eff} consistent with zero in both frozen dataset combinations and recovers $H_0 = 67.68 \pm 1.06 \text{ km s}^{-1} \text{ Mpc}^{-1}$, consistent with standard Planck Λ CDM. The ΔN_{eff} extension does not resolve the Hubble tension. We note that the ΔN_{eff} constraint reported here uses Planck PR4 / NPIPE CamSpec high- ℓ TTTEEE paired with Planck 2018 low- ℓ TT/EE + lensing (`planck_2018_lowl.EE + planck_2018_lensing.clik`); the c15 verification re-run uses the PR4-consistent counterparts `planck_2020_lollipop.lowlE` and `planckpr4lensing`, making it a release-pairing robustness rerun; the resulting 0.04σ agreement in ΔN_{eff} provides an empirical bound on pairing-induced bias at the quoted precision (see Sec. V, *Release-pairing note*). Current data neither require nor exclude a phenomenological ΔN_{eff} proxy contribution for the bounce class; CMB-S4 ($\sigma(N_{\text{eff}}) \sim 0.03$ [16]) will sharpen this constraint substantially. Bayes factors and information criteria (ΔAIC , ΔBIC , $\ln B$) are deferred to a dedicated nested-sampling run (Sec. V, *Model-comparison statistics* paragraph); we report only the parameter posteriors (ΔN_{eff} , H_0 , etc.) as load-bearing here and explicitly omit Bayes-factor or information-criterion comparisons in this Metropolis-Hastings-only analysis.

NaMaster pipeline validation.—The 500-MC pseudo- C_ℓ analysis on synthetic Λ CDM polarization skies confirms that the deconvolution pipeline recovers injected

birefringence angles with amplitude-dependent bias $\hat{\beta} - \beta_{\text{inj}} = -0.032$ to -0.040° (worst-case -0.040° at injection $\beta = 0.342^\circ$; Sec. IV) at template-fit SNR consistent with the ACT-noise floor. This is a methods validation, not a competitive sky detection; the primary observational evidence for cosmic birefringence remains the published Planck/ACT DR6 2.7–2.9 σ measurements.

Spectator-ALP consistency.—To state the verdict first: the spectator ALP *accommodates*, it does not *predict*, the observed birefringence, and the spectator-safe verdict rests on the broader $\Omega_a < 0.01$ subset (13% of the posterior mass), not on the extreme $\theta_i \leq 0.1$ sliver (0.33% mass, indicative only). Within that $\Omega_a < 0.01$ spectator-safe subset (13% of the posterior mass; median $m \simeq 40.5 H_0$, $\Omega_a < 0.01$, ESS = 461), an ALP with $f_a \sim M_{\text{Pl}}$ and posterior median $m \simeq 40.5 H_0$ (not $m \sim H_0$; the scan-prior envelope spans $m \sim H_0$ but the posterior-supported fixed- $C_{a\gamma} = 8$ fit shifts well above the prior midpoint) is consistent with the published 3.6σ joint signal, with the caveat that this spectator-safe subset ($\theta_i \sim 0.1$, a $\sim 25\times$ fine-tuning relative to the natural midpoint; fn. 6) satisfies spectator status ($C_{a\gamma} \Delta\phi/f_a \approx 10.3$ for $\beta = 0.342^\circ$, with $\Delta\phi/f_a \in [0.064, 1.19]$ over the natural box giving required $C_{a\gamma}$ from ≈ 8.6 up to ≈ 160 at the smallest displacements, and ≈ 8 – 10 in the posterior-preferred saturated-displacement regime; see §VI for the explicit numerical derivation). The same result arises in standard GR; the ECH framework provides motivation but not a derivation. Under forecast foreground-cleaning and calibration assumptions [32], a $\beta \simeq 0.27^\circ$ signal would be detected relative to $\beta = 0$ at $\sim 9\sigma$ by LiteBIRD in the early 2030s.

Data and Code Availability

All materials are at:

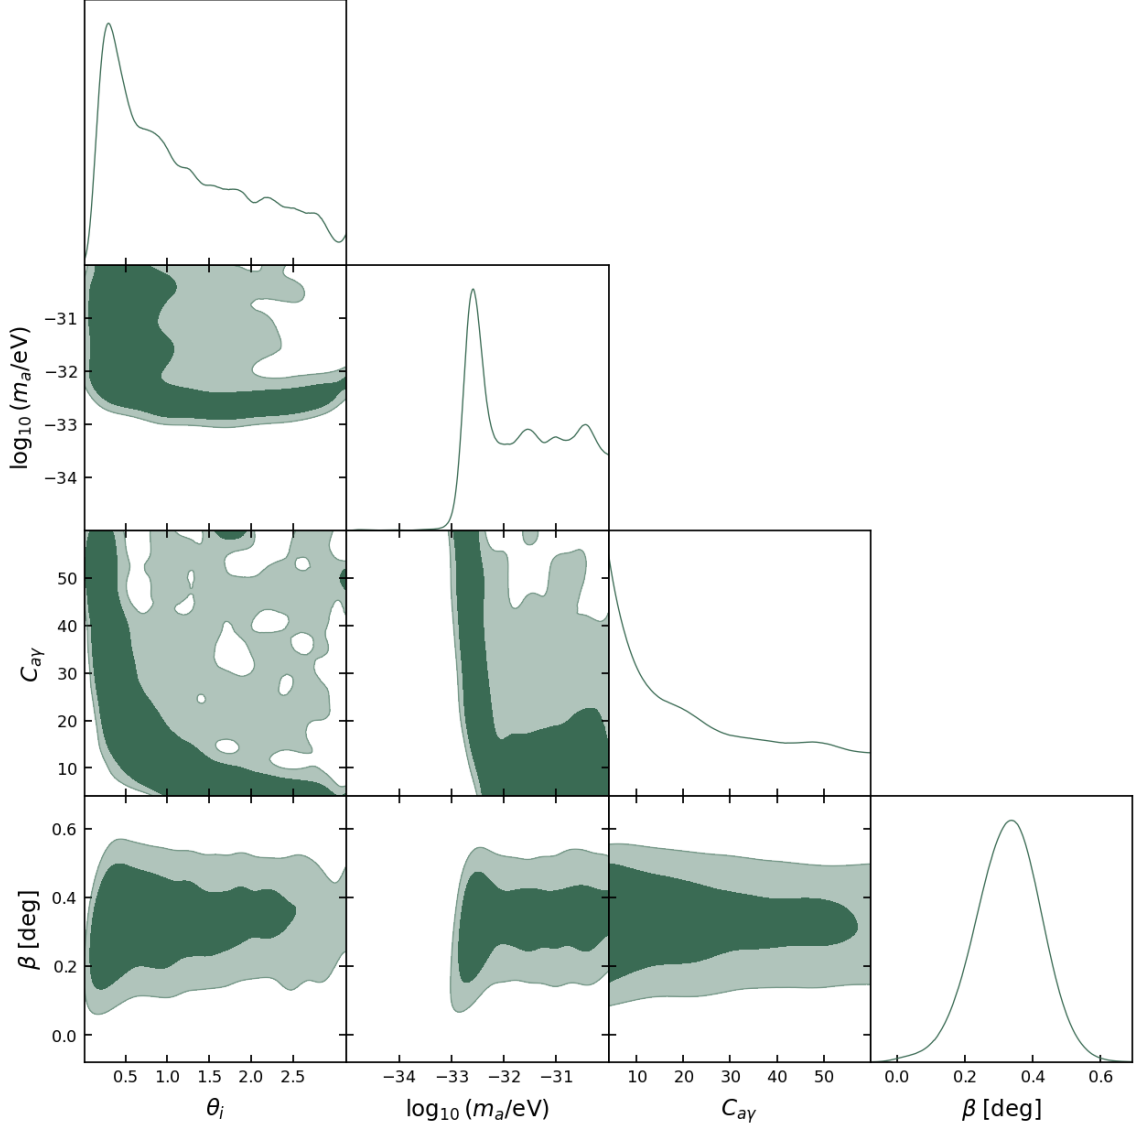


FIG. 4. Spectator-ALP joint posterior triangle from the continuous-prior cross-check configuration in which the photon anomaly coefficient is sampled freely (flat priors $C_{a\gamma} \in [4, 60]$ — shifted and extended from the earlier $[1, 30]$ to cover the posterior-supported coupling band (median 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 $H_0 = 67.7$ km/s/Mpc = 1.44×10^{-33} eV), broadening the fixed- $C_{a\gamma} = 8$ and $C_{a\gamma} \in [1, 30]$ configurations of Appendix C; 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. 6).

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

The version of the repository corresponding to this paper is identified by the in-tex v1B.0.101

stamp (defined in `paper1b_mcmc_companion.tex` `\newcommand{\paperVersion}`) and by the matching version-stamp commit in the repository git log (current snapshot commit: `b22f8cc9`); the per-version changelog (with the corresponding commit SHAs) is in

the repository root `CHANGELOG.md`. The repository is program-wide: it includes the four Cobaya YAML configurations (one per dataset combination, stock CAMB, no custom modifications), NaMaster driver scripts, and the implementation map (`IMPLEMENTATION_MAP.md`), alongside galaxy-spin pipeline code used by Paper IV [8] but not by any analysis of this paper. Frozen MCMC chains are committed (HuggingFace datasets listed in Appendix A); fresh re-verification chains can be regenerated via `reproduce_cosmology.sh` ($\sim 4\text{--}12$ h per configuration on 4 CPU cores). HuggingFace dataset URLs for the three accompanying datasets (MCMC chain diagnostics, NaMaster pipeline artifacts, ALP parameter chains) are listed below in Appendix A under “HuggingFace datasets”; DOI assignment is pending (identifiers will be inserted at submission). The URLs are also recorded in the repository `CHANGELOG.md` under the entry for `v1B.0.101`; links are preserved against future README changes via the version-stamp commit identified by the changelog. *Column-permutation warning for JSON artifacts.*—The corrected diagnostic file `parameter_summary_CORRECTED.json` in each frozen-chains directory is the *only* parameter-summary artifact served, and its values match Table I. The earlier defective export `parameter_summary.json`, which carried an off-by-one column-index bug, has been *deleted (expunged) from every public repository directory*: it is no longer present in the committed tree, so a downstream user cannot accidentally read the permuted values. Full details of the bug, the corrected file, and the column-by-column mapping are retained (for provenance) in the repository `CHANGELOG.md` (`v1B.0.59` archive entry) and in the adjacent `parameter_summary_units_README.md`; this artifact-level process record is kept in the repository, but the defective data file itself has been removed.

ACKNOWLEDGMENTS

We thank the Planck, ACT, LiteBIRD, and Cobaya collaborations for providing the data, code, and observational infrastructure used in these analyses. The author acknowledges the use of Claude (Anthropic) as an AI research assistant during systematic analysis and manuscript preparation. All scientific claims, derivations, numerical results, and bibliographic attributions were independently verified by the author. No external funding was received. Computational resources were self-funded (RunPod H200 instances).

Appendix A: 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 [33] 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 Λ CDM+ ΔN_{eff} chain directories are committed in-repo and contain the chains and diagnostics that back Table II, Table III, and the relevant Table V claim-classification entries of this paper; the ALP c_5 continuous chain backs Table IV 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 Λ CDM+ ΔN_{eff} proxy chains for independent re-verification are NOT bundled and must be regenerated locally via `reproduce_cosmology.sh` ($\sim 4\text{--}12$ h per config on 4 CPU cores). All Table II 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). The URLs are also recorded in the repository `CHANGELOG.md` under the entry for `v1B.0.101` 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/plb-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>.

Appendix B: Claims Classification

Table V 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 A.

Appendix C: ALP-MCMC Sampled Parameters, Priors, and Likelihood Stack

The ALP-MCMC results quoted in Sec. VI ($\beta_{\text{ALP}} = 0.336^\circ \pm 0.10^\circ$ 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 $[0.01, \pi]$.⁷
- $\log_{10}(m_a/\text{eV})$: 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. 4).—Three sampled 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. VI; 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 $[0.01, \pi]$.
- $\log_{10}(m_a/\text{eV})$: 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	c5_continuous	8,955	796
$\log_{10}(m_a/\text{eV})$	c5_continuous	8,955	1,114
$C_{a\gamma}$	c5_continuous	8,955	814
β_{deg} (derived)	c5_continuous	8,955	2,860
β_{free}	run3_baseline	720	265

The `c5_continuous` ESS values (~ 800 – 2860) are adequate for posterior characterisation. The `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 $\beta_{\text{obs}} = 0.342^\circ \pm 0.094^\circ$ [5] (fn. 3; 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

⁷ Backreaction / spectator-status disclosure: see fn. 6 in Sec. VI. 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 V. Claims classification for this companion paper. 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 `v1B.0.101` 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.	[5]	Eskilt et al.
Stock CAMB proxy \neq ECH theory module	Scope	Definition	§III
ALP birefringence not distinctive ECH prediction	Scope	Definition	§VI

$\Omega_a < 0.01$ 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 $\beta = 0.215^\circ \pm 0.074^\circ$ [4] is an independent cross-check quoted in Sec. VI 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] H. Golden, Structural Closure of Einstein–Cartan–Holst Dark Energy: Perturbation Transparency, Inflation- f_{NL} Tension, and Surviving Matter-Bounce Tests, (2026), companion paper, posted concurrently on arXiv [arXiv:XXXX.XXXXX — ID inserted at coordinated submission].
- [2] E. Hivon, K. M. Górski, C. B. Netterfield, B. P. Crill, S. Prunet, and F. Hansen, MASTER of the Cosmic Microwave Background Anisotropy Power Spectrum: A Fast Method for Statistical Analysis of Large and Complex Cosmic Microwave Background Data Sets, *Astrophys. J.* **567**, 2 (2002), arXiv:astro-ph/0105302 [astro-ph].
- [3] 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 $\beta = 0.30 \pm 0.11$ deg from Planck NPIPE (PR4), arXiv:2201.07682 [astro-ph.CO].
- [4] 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].
- [5] 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].
- [6] H. Golden, $f_{\text{NL}} = -35/8$ Forecast: SPHEREx Discrimination of Bounce vs. Inflation, (2026), companion paper, posted concurrently on arXiv.
- [7] 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.
- [8] H. Golden, Galaxy Chirality at Scale: 8.47M Galaxies Classified, Hemisphere Null at $p_{\text{LEE}} < 10^{-4}$, (2026), companion paper, posted concurrently on arXiv.
- [9] 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].
- [10] 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).

- [11] N. J. Popławski, Four-fermion interaction from torsion as dark energy, *General Relativity and Gravitation* **44**, 491 (2012), arXiv:1102.5667 [gr-qc].
- [12] 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].
- [13] 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].
- [14] M. P. Unger and N. J. Popławski, Big bounce and closed universe from spin and torsion, *The Astrophysical Journal* **870**, 78 (2019), arXiv:1808.08327 [gr-qc].
- [15] J. I. Kapusta and C. Gale, *Finite-Temperature Field Theory: Principles and Applications*, 2nd ed., Cambridge Monographs on Mathematical Physics (Cambridge University Press, 2006).
- [16] CMB-S4 Collaboration, K. Abazajian, *et al.*, CMB-S4 Science Book, First Edition, arXiv e-prints (2016), arXiv:1610.02743 [astro-ph.CO].
- [17] M. Brüggen, Effects of a torsion field on Big Bang nucleosynthesis, *General Relativity and Gravitation* **31**, 1935 (1999), arXiv:astro-ph/9906403 [astro-ph].
- [18] Y.-F. Cai, W. Xue, R. Brandenberger, and X. Zhang, Non-gaussianity in a matter bounce, *JCAP* **0905**, 011, arXiv:0903.0631.
- [19] 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].
- [20] 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].
- [21] D. Brout *et al.*, The Pantheon+ analysis: Cosmological constraints, *Astrophys. J.* **938**, 110 (2022), arXiv:2202.04077 [astro-ph.CO].
- [22] 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].
- [23] 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].
- [24] Planck Collaboration, N. Aghanim, *et al.*, Planck 2018 results. VI. cosmological parameters, *Astronomy & Astrophysics* **641**, A6 (2020), arXiv:1807.06209 [astro-ph.CO].
- [25] 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].
- [26] 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].
- [27] 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].
- [28] 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].
- [29] 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].
- [30] 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].
- [31] 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].
- [32] 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].
- [33] 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].