Reproducible Cosmological Proxy and Pipeline Checks: Stock-CAMB ΛCDM+$\Delta N_{\rm eff}$ MCMC, Synthetic NaMaster Recovery, and a Generic Spectator-ALP Birefringence Fit

Houston Golden[1,*]

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

(Dated: July 15, 2026, 02:53 PT | v1B.0.109)

We present three self-contained computational studies and state their limited interpretation at the outset. None implements a torsion-modified Boltzmann solver, none measures torsion, and none is evidence for Einstein–Cartan–Holst (ECH) gravity or for a bounce cosmology.

First, stock CAMB and Cobaya are used to sample a generic ΛCDM+$\Delta N_{\rm eff}$ extension. Across 309,189 frozen samples in two converged data combinations, $\Delta N_{\rm eff}$ is consistent with zero (-0.020 ± 0.169 and $+0.058 \pm 0.179$), and the extension does not remove the residual Hubble tension. This is an extra-radiation proxy, not an ECH theory run. For scale context only, we also derive the parametric minimal-contact estimate $\rho_{\rm tor}/\rho_{\rm rad} \sim (T/M_{\rm Pl})^2$; at BBN it is of order 10^{-43}–10^{-44} up to the stated thermal and convention factors, far below the proxy sensitivity.

Second, a 500-realization synthetic-CMB NaMaster test on a synthetic mask recovers injected global polarization rotations after applying the exact NaMaster bandpower-window operator: $0.270^\circ \mapsto 0.269^\circ$ and $0.342^\circ \mapsto 0.341^\circ$, with a canonical MC-mean bias $-0.0010^\circ \pm 0.0023^\circ$ (standard error). Because the skies omit Galactic foregrounds, the test does not break the cosmic-rotation/instrument-calibration degeneracy and is not a real-sky detection or systematic floor.

Third, a generic spectator-axion-like-particle (ALP) model is fitted to a published birefringence summary likelihood. The fit shows accommodation, not prediction: at fixed $C_{a\gamma} = 8$ the posterior favors $m \gg H_0$, and only a conditional subset remains spectator-safe. A separate prior-predictive calculation finds 11.6% of fixed-coupling draws and 6.1% of broad-coupling draws within the published 1σ band; these are unconditional prior-sensitivity fractions, not the cost or probability of the spectator subset. Of the 34,947 trajectories satisfying $\Omega_a < 0.01$, the corresponding conditional counts include zero observed fixed-coupling draws (95% Wilson upper bound 0.011%) and 863 (2.469%; MC standard error 0.083 percentage points) for the broad coupling prior. These are Monte Carlo frequencies, not evidence or posterior probabilities. The same model exists in standard GR and has no ECH-specific content.

Frozen chains, configurations, synthetic-recovery outputs, and the prior-predictive script are publicly archived. The companion's contribution is reproducibility and quantitative scope control for these three studies.

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

CONTENTS

* houston@hubify.com

TABLE I. Principal symbols used in the birefringence / spectator-ALP consistency check (Sec. VI) and the ΛCDM+$\Delta N_{\rm eff}$ 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); $\beta_{\rm obs}$ published datum, $\beta_{\rm ALP}$ 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, $\rho_a(z{=}0)/\rho_{\rm crit,0}$; spectator cut $\Omega_a < 0.01$
f_a	ALP decay constant (GeV); fiducial $f_a \sim M_{\rm Pl}$
m	ALP mass, quoted in units of H_0; scan-prior region $m \sim H_0$
$M_{\rm Pl}$	Reduced Planck mass, $(8\pi G)^{-1/2} \approx 2.44 \times 10^{18}$ GeV
$\Delta N_{\rm eff}$	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

I. INTRODUCTION

This paper reports three limited numerical studies and their reproducibility artifacts: a stock-CAMB extra-radiation proxy, a synthetic NaMaster recovery test, and a generic spectator-ALP fit. It does *not* implement, test, or verify a torsion-modified cosmology. Its relation to Paper I(a) is bibliographic and programmatic, not evidentiary. The present paper also records the dimensional scale estimate $\rho_{\rm tor}/\rho_{\rm rad} \sim (T/M_{\rm Pl})^2$ for the minimal algebraic-torsion contact operator; because the thermal prefactor is only parametric, we use it as scale context rather than as a precision prediction. Paper I(a) derives the minimal algebraic-torsion contact interaction, evaluates that specific channel in a standard mean-field NJL model, and records the classical transparency of the constant-Immirzi Holst term on the canonical scalar torsion-free branch. It does not claim a complete gravitational-EFT no-go or a derived ECH dark-energy observable. The present paper is logically separate: its three numerical analyses are generic proxy, pipeline, and consistency studies. They neither support nor test the action-level claims of Paper I(a), and Paper I(a) does not supply an ECH interpretation for them. The two papers are coordinated concurrent submissions, but the three 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.

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. *ΛCDM+$\Delta N_{\rm eff}$ MCMC proxy* (Secs. II and III): stock CAMB Boltzmann solver with $\Delta N_{\rm eff}$ 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 ΛCDM polarization skies. **Not a competitive sky detection.** The diagonal-covariance template score (e.g., 20.01) is a heuristic recovery diagnostic for injected MC signals and omits inter-bin covariance; it is not the significance of the CMB sky measurement, which remains the published 2.7–2.9σ 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 $m \gg H_0$ (weighted median $37.2\,H_0$), well above the prior midpoint. A spectator-safe interpretation is reported conditionally at $\Omega_a < 0.01$; that subset has median $\theta_i = 0.211$ and is distinct from the indicative $\theta_i \leq 0.1$ sliver (Sec. VI). **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.—We do not derive an ECH dark-energy, birefringence, or bounce observable, and we do not use the numerical results below to support an ECH claim. Paper I(a) [1] contains the separate minimal-action contact-channel and classical-transparency analysis.

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

We extend ΛCDM by $\Delta N_{\rm eff}$ as a generic phenomenological extra-radiation 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 [2] 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 $\Delta N_{\rm eff}$ in stock CAMB; in this configuration the $\Delta N_{\rm eff}$ posterior is consistent with zero and the H_0 posterior is dominated by Planck. This generic extra-radiation extension does not resolve cosmological tensions at the present data precision.

III. STOCK-CAMB ΛCDM+$\Delta N_{\rm eff}$ MCMC: GENERIC RADIATION-PROXY TEST (NOT A SPIN-TORSION THEORY MODULE)

Scope statement.—"MCMC verification" refers throughout to a stock CAMB run of ΛCDM extended by $\Delta N_{\rm eff}$ 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 with no theory-specific identification. It does *not* verify the spin-torsion theory module itself; that would require a derived stress tensor and a bespoke modified Boltzmann implementation. The dimensional estimate in Sec. III A is shown only to compare scales and is not used to reinterpret the stock-CAMB posterior as an ECH constraint.

A. Parametric radiation-era scale of the minimal contact operator

For scale context, we estimate the parametric contribution of the minimal Einstein–Cartan–Holst (ECH) contact operator to the early-universe energy budget. 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 [3]. Integrating torsion out of the action leaves an effective four-fermion contact interaction of Nambu–Jona-Lasinio type [4, 5],

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

where γ is the Barbero–Immirzi parameter of the Holst term [5, 6]; in the pure Einstein–Cartan–Sciama–Kibble limit $\gamma\to\infty$ the bracket $\to 1$ and Eq. (1) reduces to the canonical $-\frac{3}{2}\pi G_N(\bar\psi\gamma^5\gamma^\mu\psi)^2$ [4]. The coefficient follows from the declared minimal action; the thermal expectation value below is nevertheless an order-of-magnitude estimate rather than a precision finite-temperature calculation.

The crucial feature of Eq. (1) is that the coupling is $\propto \kappa^2 = 8\pi G_N \sim M_{\rm Pl}^{-2}$: the interaction is Planck-suppressed. In a thermal fermion gas the parametric contact energy density is $\rho_{\rm tor} \sim \alpha\, n_f^2$ with $\alpha = \kappa^2(\hbar c)^2/32$ and number density $n_f \propto T^3$ for relativistic species [7]. Explicitly, the energy density contributed by the dimension-6 contact operator of Eq. (1) is the thermal expectation value $\rho_{\rm tor} \sim \kappa^2\langle(\bar\psi\gamma^5\gamma^\mu\psi)^2\rangle_T$; 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 [8], with $\mathcal{O}(1)$ coefficients dropped rather than a rigorous NJL finite-T computation; *a full finite-temperature NJL computation of the exact prefactor is outside the scope of this companion; the omitted prefactor prevents identifying this ratio with an effective-neutrino count*), so that

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

since $\rho_{\rm rad} \sim T^4$. 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 [7]. For scale comparison, the same parametric ratio at the epochs probed by the data is

$$\frac{\rho_{\rm tor}}{\rho_{\rm rad}} \sim \left(\frac{T}{M_{\rm Pl}}\right)^2 \sim \begin{cases} 1.7\times10^{-43}, & T\sim 1\ \text{MeV},\\ 1.1\times10^{-56}, & T\sim 0.26\ \text{eV},\end{cases} \tag{3}$$

at BBN and recombination respectively, evaluated with the reduced Planck mass $M_{\rm Pl} = (8\pi G_N)^{-1/2} \simeq 2.44\times 10^{18}$ GeV used throughout (and in Eq. 2), so the boxed numbers are reproducible directly from $(T/M_{\rm Pl})^2$. These numbers estimate $\rho_{\rm tor}/\rho_{\rm rad}$, not $\Delta N_{\rm eff}$: the latter conversion would require a specified thermal spin ensemble and its numerical prefactor. Within the parametric treatment just stated, the minimal contact operator gives a negligible radiation-era scale. 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 [9]. We therefore do *not* claim a specific non-zero torsion $\Delta N_{\rm eff}$; Eq. (3) reports only the leading parametric energy-density ratio.

This estimate does not reframe the MCMC as a theory test. The stock-CAMB ΛCDM+$\Delta N_{\rm eff}$ run remains a generic extra-radiation proxy. Its data-driven 95% bound $\Delta N_{\rm eff} \lesssim 0.31$–$0.40$ (Sec. V) is simply many orders above the contact-operator scale estimate; agreement between two null-sized quantities is not evidence for the underlying theory. Near Planckian temperatures the contact-EFT and thermal approximations would themselves require a UV completion, a regime outside both a stock Boltzmann code and this paper.

The proxy run (Cobaya v3.6.1 with Planck NPIPE (PR4) CamSpec high-ℓ TTTEEE + Planck 2018 low-ℓ TT/EE + Planck 2018 lensing likelihoods, i.e. `planck_NPIPE_highl_CamSpec.TTTEEE` + `planck_2018_lowl.TT` + `planck_2018_lowl.EE` + `planck_2018_lensing.clik` in the Cobaya YAML) has produced two frozen dataset combinations with publication-quality convergence (limitation: the PR4/NPIPE high-ℓ + 2018-release low-ℓ/lensing mixture is the standard Cobaya pairing; the c15 verification re-run uses `planck_2020_lollipop.lowlE` and `planckpr4lensing` in place of the frozen-chain `planck_2018_lowl.EE` and `planck_2018_lensing.clik`, making c15 a release-pairing sensitivity rerun; because it does not meet the stated convergence threshold, its 0.04σ central-value agreement is corroborative only and does not bound pairing-induced bias; 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 ΛCDM+$\Delta N_{\rm eff}$ null-consistency test: the data are consistent with $\Delta N_{\rm eff} = 0$ in stock CAMB. Sampling configuration (per the on-disk Cobaya YAMLs): $\Delta N_{\rm eff}$ enters as `nnu` with a flat prior $N_{\rm eff} \in [2.046, 5.046]$ (i.e. $\Delta N_{\rm eff} \in [-1, +2]$); the neutrino mass is held at the CAMB default $\Sigma m_\nu = 0.06\,$eV with one massive eigenstate (`num_massive_neutrinos`=1), and $Y_{\rm He}$ 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): $Y_{\rm He}$ is set self-consistently with $N_{\rm eff}$ at each MCMC sample. The prior range $N_{\rm eff} \in [2.046, 5.046]$ remains within the calibrated domain of the CAMB BBN module; no free-$Y_{\rm He}$ control run was performed, as the default BBN-consistent track is the standard choice for this type of proxy analysis. The prior allows negative $\Delta N_{\rm eff}$ (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 $\Delta N_{\rm eff} < 0$; 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 $\Delta N_{\rm eff} \geq 0$. Under that physically motivated restriction $\Delta N_{\rm eff} \geq 0$, the one-sided 95% upper limit from the full-tension chain is $\Delta N_{\rm eff} < 0.31$, defined precisely as the 95th percentile of the posterior *renormalised on* $\Delta N_{\rm eff} \geq 0$ (we discard the $\Delta N_{\rm eff} < 0$ tail of weighted samples and rescale the surviving weights so $\int_0^\infty p(\Delta N_{\rm eff})\, d\Delta N_{\rm eff} = 1$, 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 $\Delta N_{\rm eff} < 0.40$ 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 $\Delta N_{\rm eff} \geq 0$, 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 $\Delta N_{\rm eff}$ proxy: generic extra radiation, not a theory discriminator. We emphasize that

[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, $176{,}240 \times 0.7 + 132{,}949 \times 0.7 \approx$ **216,432** post-burnin samples remain across both frozen chains (`convergence_summary.json`). For the full-tension subset specifically, $176{,}240 \times 0.7 \approx$ **123,368** post-burnin (the **119,617** figure in Fig. 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.

the stock-CAMB ΛCDM+$\Delta N_{\rm eff}$ 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 [3, 6] is dimension-6 and $M_{\rm Pl}^{-2}$-suppressed[2]: its leading Boltzmann effect is a scattering-amplitude shift, not a relativistic species, and it does *not* produce a $\Delta N_{\rm eff}$ at recombination. The proxy run finds $\Delta N_{\rm eff} = -0.020 \pm 0.169$ (full-tension) and $+0.058 \pm 0.179$ (Planck+BAO+SN), both consistent with zero. This constrains only the stock-CAMB extra-radiation parameter. It neither confirms nor discriminates an ECH or bounce model, because no such model is implemented in the likelihood calculation.

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 $M_B - 5\log_{10}(h) \approx$ const, with $h = H_0/(100\,{\rm km\,s^{-1}\,Mpc^{-1}})$ (equivalently $M_B - 5\log_{10}(H_0/[{\rm km\,s^{-1}\,Mpc^{-1}}]) + 10 =$ const), along the SN distance-modulus degeneracy. At the SH0ES [2] 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 $\Delta N_{\rm eff}$ 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`).

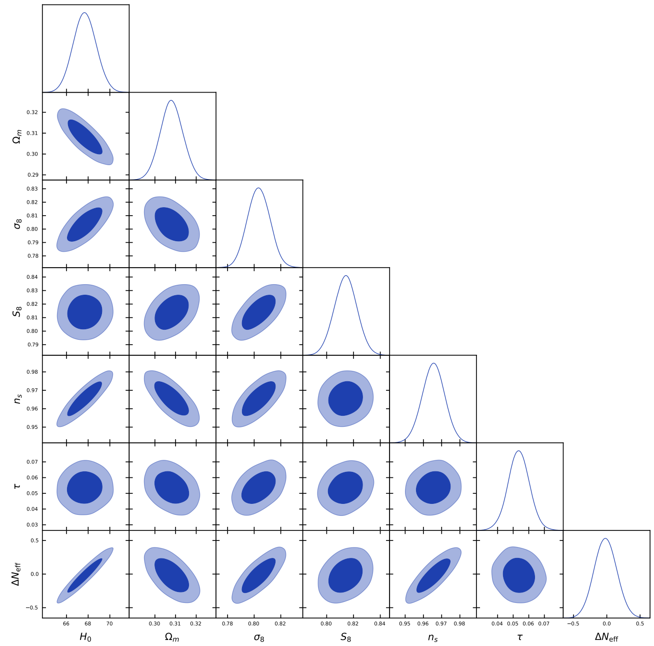

FIG. 1. Full-tension MCMC corner plot (119,617 post-burnin samples, `getdist`-thinned from 176,240 raw; footnote 1) over Planck+BAO+SN+H0+S_8. The $\Delta N_{\rm eff}$ posterior is consistent with zero (-0.020 ± 0.169), confirming no additional relativistic species at recombination.

Key finding.—Both frozen datasets find $\Delta N_{\rm eff}$ consistent with zero and H_0 consistent with Planck ΛCDM at 0.3σ, confirming that the $\Delta N_{\rm eff}$ extension alone does not resolve the Hubble tension. Current data neither require nor exclude a small positive phenomenological $\Delta N_{\rm eff}$ extra-radiation contribution; CMB-S4 ($\sigma(N_{\rm eff}) \sim 0.03$ [10]) will sharpen this constraint substantially relative to the current Planck+BAO+SN level.

IV. DATA METHODS: CMB E-B ANALYSIS

Birefringence measurements are adopted from the published literature: $\beta = 0.30° \pm 0.11°$ (Planck NPIPE [11]) and $\beta = 0.215° \pm 0.074°$ (ACT DR6 [12]). The spectator ALP analysis (Sec. VI) instead uses the Eskilt–Komatsu joint WMAP+Planck summary likelihood ($\beta =$

[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 $\Lambda_{\rm strong} \sim M_{\rm Pl}/\sqrt{\gamma_{\rm BI}}$ set by the inverse Barbero–Immirzi parameter $\gamma_{\rm BI}$ (cf. [6], where the contact interaction is derived; the strong-coupling scale itself is a heuristic EFT-validity estimate); above this scale the contact-operator description ceases to be controlled and a specified ultraviolet completion with additional propagating degrees of freedom would be required. Minimal Holst gravity itself does not supply propagating torsion: its torsion remains an algebraic auxiliary field. All recombination-era and late-time observables addressed in this paper sit at energies $E \ll \Lambda_{\rm strong}$, where the EFT treatment is controlled.

TABLE II. ΛCDM+$\Delta N_{\rm eff}$ 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 $\Delta N_{\rm eff}$ 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 $\Delta N_{\rm eff} \geq 0$ (posterior renormalised on the non-negative half and 95th percentile read from the truncated CDF), the one-sided 95% upper limits are: full-tension $\Delta N_{\rm eff} < 0.31$; Planck+BAO+SN $\Delta N_{\rm eff} < 0.40$ (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 $\Delta S_8 = 5.5 \times 10^{-5}$ trapezoidal grid with 4,001 points spanning $S_8 \in [0.70, 0.92]$); 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
$\Delta N_{\rm eff}$	-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]	0.000974	0.000971
Min integrated-autocorrelation ESS	4,761.37 ($\Delta N_{\rm eff}$)	4,692 (σ_8)

[a] The archived Planck+BAO+SN $\Delta N_{\rm eff}$ row has $\hat{R} - 1 = 9.71 \times 10^{-4}$; 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 ($A_{\rm planck}$, amp$_{143}$, amp$_{217}$, amp$_{143\times217}$, n_{143}, n_{217}, $n_{143\times217}$, calTE, calEE) + 1 supernova absolute magnitude (M_b, entering via the Pantheon+ SN and SH0ES `H0.riess2020Mb` likelihoods rather than the Planck likelihood stack). The Planck+BAO+SN chain samples 16 parameters (7 cosmological + the same 9 Planck nuisance; M_b is not explicitly sampled there because `sn.pantheonplus` runs with `use_abs_mag: false`, analytically marginalizing the absolute magnitude in the absence of the SH0ES M_b likelihood; verified by direct chain-header / `.updated.yaml` blocking-list inspection). The ESS row uses the integrated-autocorrelation estimator in the frozen diagnostics: `reproducibility/cosmology/frozen/full_tension_20260311_1728/diagnostics/freeze_diagnostics_CORRECTED.json` reports 5,089.94 (H_0), 4,761.37 ($\Delta N_{\rm eff}$), and 6,535.81 (τ), hence the reported minimum 4,761.37 at $\Delta N_{\rm eff}$; `reproducibility/cosmology/frozen/planck_bao_sn_20260312_1954/diagnostics/convergence_report.txt` reports a minimum 4,692 at σ_8. The row-expanded `reproducibility/cosmology/convergence_latest.csv` is retained as a separate monitoring export with a weight/count-based ESS convention of order 3×10^5 and is not the source of the Table I ESS values.

$0.342^\circ \pm 0.094^\circ$ [13])[3] as its primary constraint; the Planck NPIPE and ACT DR6 values above are quoted only for context, without combining their differently constructed summaries into a new significance. (The quoted 3.6σ is the ratio $\beta/\sigma = 0.342/0.094 \approx 3.6$; our likelihood is a Gaussian summary in the published mean and width, not a refit of the Eskilt–Komatsu posterior shape.)

Scope note.—*The NaMaster pseudo-C_ℓ analysis below validates only the algebraic $E \to B$ mode-coupling deconvolution on foreground-free synthetic skies; it does **not** establish a real-sky systematic budget (the β–α calibration degeneracy, galactic foregrounds, instrumental beams, and anisotropic noise are all absent by construction and explicitly out of scope). The high diagonal-covariance template scores (e.g.,* 20.01, 25.32*; fn. 4) are heuristic recovery diagnostics that omit inter-bin co-*

[3] Eskilt & Komatsu 2022 disambiguation: the published PRD paper [13] (PRD 106:063503, arXiv:2205.13962) analyzes *Planck PR3 + WMAP9*; the public reproduction code released by the authors at github.com/LilleJohs/Cosmic_Birefringence was subsequently updated to use *Planck PR4 / NPIPE*. Throughout this paper, the labels "PR4/NPIPE" attached to the Eskilt+Komatsu likelihoods refer to the code-repository dataset; the ALP MCMC uses only the scalar Gaussian summary likelihood $\beta = 0.342^\circ \pm 0.094^\circ$ and does not depend on the PR3/PR4 map-level distinction except through the provenance of that summary value. The abstract $\beta = 0.342^\circ \pm 0.094^\circ$ (3.6σ) headline is from the published PR3+WMAP9 joint analysis. The repository README is the authoritative source for the dataset attribution in the executed pipeline.

TABLE III. Exact-bandpower-window NaMaster recovery checks on foreground-free synthetic skies ($N_{\rm side} = 512$, $N = 500$ per row; deterministic seeds 42–541). Where reported, the uncertainty is the standard error (SE) of the Monte Carlo mean. These are pipeline checks, not real-sky systematic bounds.

Configuration	$\beta_{\rm inj}$	$\hat{\beta}$	Signed bias	SE
Canonical native-coordinate window, $f_{\rm sky} = 0.32$	$+0.270°$	$+0.269°$	$-0.001°$	$0.002299°$
Negative sign, $f_{\rm sky} = 0.32$	$-0.270°$	$-0.270°$	$0.000°$	$0.002200°$
Large sky, $f_{\rm sky} = 0.85$	$+0.270°$	$+0.270°$	$0.000°$	$0.001412°$
Intermediate sky, $f_{\rm sky} = 0.65$	$+0.270°$	$+0.268°$	$-0.002°$	$0.001582°$
Apodization $0.5°$ / $3.0°$	$+0.270°$	$+0.269°$ / $+0.269°$	$-0.001°$ / $-0.001°$	—
Native-latitude cut $\lvert\text{lat}\rvert > 30°$	$+0.270°$	$+0.269°$	$-0.001°$	—
CAMB-lensed BB injection	$+0.270°$	$+0.270°$	$0.000°$	—
B-purified field	$+0.270°$	$+0.269°$	$-0.001°$	—

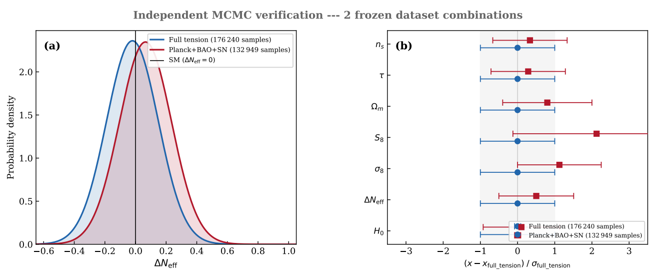

FIG. 2. $\Delta N_{\rm eff}$ marginal posterior comparison across the two frozen dataset combinations of Table II (176,240 and 132,949 samples). Panel (a): Gaussian summaries of the $\Delta N_{\rm eff}$ marginal posteriors at the Table II means $\pm 1\sigma$, with the Standard-Model value $\Delta N_{\rm eff} = 0$ marked. Panel (b): all seven Table II parameters, normalized to the full-tension mean and σ. Both combinations recover $\Delta N_{\rm eff}$ consistent with zero at $< 1\sigma$; the headline "full-tension" stack (Planck+BAO+SN+H$_0$+S_8) gives $\Delta N_{\rm eff} = -0.020 \pm 0.169$. No evidence for a recombination-era $\Delta N_{\rm eff}$ 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 $\Delta N_{\rm eff}$ in stock-CAMB (see § III scope note).

variance and must not be conflated with the published Planck/ACT DR6 2.7–2.9σ *sky detection. The* β–α *degeneracy that dominates real-sky birefringence measurements is broken only by unrotated galactic foregrounds, which are absent from these synthetic skies by construction (cf. abstract).*

Pipeline configuration.—The pseudo-C_ℓ analysis follows the NaMaster framework [14]. *Simulated skies.*—Each Monte Carlo realization is a synthetic ΛCDM CMB Q/U sky generated directly at $N_{\rm side} = 512$ (`healpy.synfast`, harmonic content to $\ell_{\rm max} = 2N_{\rm side} = 1024$) 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 $N_{\rm side} = 512$ 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.—The canonical check uses a synthetic HEALPix native-coordinate latitude window. The same native latitude returned by `healpy.pix2ang(..., lonlat=True)` is required to satisfy both $\lvert\text{lat}\rvert > 20°$ and $-65° < \text{lat} < +25°$; no coordinate rotation is applied. The binary window is Gaussian-smoothed at 2° FWHM, giving $f_{\rm sky} = 0.32260$ (analytic unapodized area 0.32244). It is not a Galactic-plus-equatorial mask and does not represent the ACT or any other survey footprint. No E/B purification is applied (`purify_b=False`); residual $E \to B$ leakage from the apodized mask is absorbed into the measured pipeline-recovery bias, which is the quantity this validation is designed to characterize.

Mode-coupling matrix and binning.—The $M_{\ell\ell'}$ matrix is computed via `NmtWorkspace.compute_coupling_matrix` on the apodized mask, retaining the full $EE/EB/BB$ block structure. (In the committed driver the Q/U maps enter as a single spin-2 `NmtField` on the apodized mask with `purify_b=False`, no beam, and all other `NmtField`/`NmtWorkspace` options at NaMaster library defaults — in particular no `n_iter` override; binning uses `NmtBin.from_edges` with 20 linear integer-edge bins from `np.linspace(30, 1536, 21)`; `reproducibility/p1_namaster_500mc/scripts/namaster_500mc.py`.) Spectra are band-power-binned into 20 linear bins spanning $30 \le \ell \le 3N_{\rm side} = 1536$ (bins above the map band limit $\ell = 1024$ carry noise only). The recovery model uses the workspace's exact four-spectrum bandpower-window tensor `NmtWorkspace.get_bandpower_windows`: the rotated $[EE, EB, BE, BB]$ theory is contracted through the same window operator as the decoupled estimator. Before every ensemble, the code verifies this contraction against `decouple_cell(couple_cell(theory))` to 10^{-10} maximum absolute tolerance. Thus no effective-ℓ

bin-centre approximation enters the current results. The previous effective-ℓ outputs are preserved only as explicitly superseded historical artifacts.

Noise model and injections.—Each realization adds white noise at the ACT-like level $\Delta_P = 10\,\mu\mathrm{K}\cdot\mathrm{arcmin}$ (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 $\sigma_{\rm pix} = \Delta_P/\sqrt{\Omega_{\rm pix}}$ with $\Omega_{\rm pix}$ the $N_{\rm side} = 512$ pixel area expressed in arcmin^2 ($\Omega_{\rm pix} = 47.21\,\mathrm{arcmin}^2$, giving $\sigma_{\rm pix} = 10/\sqrt{47.21} = 1.455\,\mu\mathrm{K}$; algebraically identical to the standard $\sigma_{\rm pix} = \Delta_P\,[\pi/(180{\times}60)]/\sqrt{\Omega^{(\rm sr)}_{\rm pix}}$; Δ_P is in thermodynamic CMB units, $\mu\mathrm{K_{CMB}}{\cdot}\mathrm{arcmin}$, matching the `synfast` $\mu\mathrm{K_{CMB}}$ maps — bandpass / colour-correction conversions do not enter the synthetic pipeline), and draws independent Gaussian realizations with the *same* $\sigma_{\rm pix}$ for Q and U (no $\sqrt{2}$ factor; `reproducibility/p1_namaster_500mc/scripts/namaster_500mc.py`). The $\beta = 0.27°$, $\beta = 0.342°$, 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_b^{EB} bandpowers to the exact windowed rotated theory on a β grid (a uniform 2,001-point grid of angles spanning $\beta \in [-1°, +1°]$ at 0.001° step, with β converted to radians when evaluating sin and cos; `namaster\_500mc.py`). Explicitly, the estimator minimises

$$\chi^2(\beta) = \sum_b \left[C_b^{EB,\rm decoupled} - C_b^{EB,\rm win}(\beta)\right]^2, \qquad (4)$$

where, writing $j \in \{EE, EB, BE, BB\}$,

$$C_b^{EB,\rm win}(\beta) = \sum_j \sum_\ell W_{b\ell}^{EB,j}\, C_\ell^j(\beta), \qquad (5)$$

$$C_\ell^{EB}(\beta) = \frac{\sin(4\beta)}{2}\left(C_\ell^{EE} - C_\ell^{BB}\right). \qquad (6)$$

and the corresponding rotated EE, BE, and BB spectra are included in the same tensor contraction. The sum in Eq. (4) runs over all 20 bins. It is deliberately unweighted—there is no σ_b^2 divisor—while an inverse-variance-weighted fit is reported as a robustness cross-check. Spectra above the simulated map band limit are zero-padded before window contraction; no pixel-window deconvolution or effective-ℓ evaluation is applied. The production suite used non-negative injections; a dedicated 500-realization $\beta = -0.27°$ rerun at $f_{\rm sky} = 0.32$ (`reproducibility/p1_namaster_500mc/results/exact_window_500mc/shards/declared_negative_beta_fsky_0p32.json`) recovers $-0.270°$ with SE 0.002200°, consistent with zero signed bias and confirming sign-symmetric recovery.

Reproducibility.—The exact-window driver, deterministic seeds (`seed_base=42`), output summary, per-realization bandpowers, and robustness batteries are archived under `reproducibility/p1_namaster_500mc/results/exact_window_500mc/`. The April pod output mirrored at `pipelines/h200_results/pod1_namaster_umap_2026-04-29/` is retained for provenance but superseded because its fit used an effective-ℓ theory template.

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

$$\hat\beta_{\rm NaMaster} = 0.269° \quad (\text{500-MC sample mean of } \hat\beta). \qquad (7)$$

This MC-recovery value 0.269° is a pipeline-validation figure on synthetic skies and is *not directly comparable* to the WMAP+Planck 3.6σ sky-detection significance of $\beta_{\rm obs} = 0.342° \pm 0.094°$ [13]: the significance figure is a sky-measurement null-rejection statistic, whereas 0.269° is the converged MC sample mean for an *injected* signal on noise-only synthetic skies. The pipeline-recovery bias, defined throughout as $\Delta\hat\beta \equiv \hat\beta - \beta_{\rm inj}$, is $-0.0010°$ for the canonical injection. The per-realization angle scatter is 0.0514°, so the standard error of the 500-MC mean is 0.0023°: the residual is statistically consistent with zero. The earlier claimed multiplicative under-recovery came from evaluating the theory at effective-ℓ bin centres and is superseded by the exact-window result; it is not carried forward as a pipeline systematic. The validation remains foreground-free and supplies *no* systematic floor for real-sky β measurements.[4] For $\beta = 0.342°$ (the published joint WMAP+Planck value [13]), the pipeline recovers 0.341° at template-fit SNR = 25.32; for $\beta = 0$, the recovered angle is $-0.001°$ with template-fit SNR 0 (null check; `reproducibility/p1_namaster_500mc/results/exact_window_500mc/summary.json`). All three recovered means are displaced by $-0.001°$ on the fixed 0.001° fit grid. The exact-window test therefore finds no resolved multiplicative under-recovery. This remains a methodology cross-check, not a competitive sky measurement.

Synthetic sky-fraction sweep.—To test sky-fraction sensitivity without representing survey geometry, we repeat the $\beta = 0.27°$ injection-recovery exercise with pure-latitude binary cuts $|b| > 8.6269°$ and $|b| > 20.4873°$, each smoothed at 2° FWHM. They yield apodized $f_{\rm sky} \approx 0.850$ and 0.650, respectively; there is no declination or point-source cut, so neither mask represents Planck, ACT, or any other survey footprint. The negative-sign

[4] The "pipeline-recovery SNR= 20.01" figure (and analogously 25.32 for the $\beta = 0.342°$ injection below) is the diagonal-covariance template score of the driver script, $\mathrm{SNR^{tmpl}} \equiv \left[\sum_b (C_b^{EB,\rm th}/\sigma_b^{\rm MC})^2\right]^{1/2}$, where $C_b^{EB,\rm th}$ is the full exact-window template relative to its null value and $\sigma_b^{\rm MC}$ is the per-bin scatter across the 500 realizations. Because inter-bin covariance is omitted, it is a heuristic diagnostic rather than a matched-template significance, and is not the significance of the recovered angle itself. The canonical per-realization angle scatter is $\sigma_\beta = 0.0514°$, whereas the MC-mean standard error is smaller by $\sqrt{N} = 22.4$.

$f_{\rm sky} \approx 0.32$ row analogously uses the pure-latitude cut $|b| > 42.8436°$. All use the same apodization recipe, noise level, and $N = 500$ MC realizations ($N_{\rm side} = 512$). The recovered values are $\hat\beta = 0.270°$ at $f_{\rm sky} = 0.85$ (SE $0.001412°$) and $\hat\beta = 0.268°$ at $f_{\rm sky} = 0.65$ (SE $0.001582°$). Their signed biases are respectively $0.000°$ and $-0.002°$, consistent with the exact-window canonical result at the precision of the fit grid. The sweep therefore provides no evidence for a resolved multiplicative recovery bias.

Robustness battery.—All completed battery rows use the same exact bandpower-window contraction as the canonical estimator ($N = 500$ and seeds 42–541 per row; shard artifacts under `reproducibility/p1_namaster_500mc/results/exact_window_500mc/shards/`). The canonical refit gives $0.269°$; changing the apodization FWHM to $0.5°$ or $3.0°$, or enlarging the synthetic native-latitude cut to $|\text{lat}| > 30°$, also gives $0.269°$. Restricting the canonical fit to $\ell \leq 1024$ gives $0.269°$, while inverse-variance weighting gives $0.270°$. Replacing the proxy BB injection with a CAMB-lensed ΛCDM BB spectrum gives $0.270°$. Enabling NaMaster B-mode purification gives $0.269°$, a signed bias of $-0.001°$, for the same $N = 500$ realizations and seeds 42–541. This row uses the same exact `NmtWorkspace.get_bandpower_windows` tensor contraction; its result artifact has SHA-256 `208d168546c73f5521e96ead1c2be211534a21f2096fac45647f6d5999f02e7f`. Thus every completed exact-window row lies within $0.002°$ of the injected value; the old effective-ℓ attribution to estimator weighting or injected-BB shape is not carried forward.

V. COSMOLOGICAL FITS AND MODEL COMPARISON

A. Datasets and Configuration

We analyze four dataset combinations: (1) Planck CMB (NPIPE (PR4) CamSpec high-ℓ TTTEEE + 2018 low-ℓ TT/EE + 2018 lensing) [15] (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/Lyα(auto and $\times$QSO) [16] + SDSS DR7 MGS [17] + 6dFGS [18]; (3) +Pantheon+ [19]; (4) +SH0ES M_B-anchor [2] + DES Y3 S_8 [20]. No DESI BAO likelihood enters the frozen ΛCDM+$\Delta N_{\rm eff}$ chains. The exact Cobaya likelihood blocks of all four chains are listed in Table IV. For scope clarity: Table II summarizes only the two *frozen* ΛCDM+$\Delta N_{\rm eff}$ chains (Planck+BAO+SN and full-tension), and the Planck-only (accumulating) and Planck+BAO (diagnostic, not in headline) rows of Table IV feed neither table. Parameter estimation uses Cobaya [21] (v3.5 original; v3.6.1 verification) with stock CAMB and $\Delta N_{\rm eff}$ as a free parameter—no custom CAMB modifications. The extended parameter space adds $\Delta N_{\rm eff}$ 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: ΛCDM+$\Delta N_{\rm eff}$ proxy

The two load-bearing posteriors from the frozen chains are $\Delta N_{\rm eff} = -0.020 \pm 0.169$ (full-tension) and $+0.058 \pm 0.179$ (Planck+BAO+SN), both consistent with zero. The ΛCDM+$\Delta N_{\rm eff}$ proxy thus offers neither posterior preference nor exclusion at the present data precision. The $\Delta N_{\rm eff}$ extension alone does not resolve the Hubble tension.

Independent re-run cross-check (this version).— A dedicated re-run of the Planck+BAO+SN ΛCDM+$\Delta N_{\rm eff}$ configuration (matching likelihood stack for the SDSS BAO block and high-ℓ CamSpec component; the c15 low-ℓ EE and lensing likelihood names, `planck_2020_lollipop.lowlE` and `planckpr4lensing`, differ from the frozen-chain names in Table IV, `planck_2018_lowl.EE` and `planck_2018_lensing.clik`; the pairing mismatch is noted as a known limitation in the *Release-pairing note* below) reached $R-1 = 0.0147$ on an independent compute pod and therefore did *not* meet this paper's $R-1 < 0.01$ convergence threshold (40,349 raw rows, 28,245 post-burn-in at 30%, 107,853 GetDist sum-of-weights — not the integrated-autocorrelation ESS) and gives $\Delta N_{\rm eff} = +0.0514 \pm 0.171$, in 0.04σ agreement with the frozen $+0.058 \pm 0.179$ quote above. It is used only as a corroborative robustness diagnostic, not as a replacement headline chain. 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 corroborated at the present diagnostic precision, while the frozen chains remain the headline result; full chain artifacts and posterior summary are committed at `reproducibility/cosmology/chains/w0wa_quintom_desi_dr2/c15_converged/`. The `c15_converged` directory name is a legacy artifact label; it does not override the numerical threshold statement above. *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.clik` (Planck 2018 releases). The c15 verification re-run instead uses `planck_2020_lollipop.lowlE` and `planckpr4lensing` for the low-ℓ EE and lensing components, making c15 a release-pairing sensitivity rerun (PR4-consistent low-ℓ/lensing) rather than an identical-likelihood rerun of the frozen chains. Its 0.04σ central-value agreement is corroborative only: the chain did not meet $R-1 < 0.01$, so it cannot bound pairing-induced bias (see Sec. III, limitation note).

TABLE IV. Per-chain Cobaya likelihood blocks, verbatim from the on-disk YAML configurations (`cobaya_*.yaml`; frozen-chain `spin_torsion.input.yaml`). All four ΛCDM+$\Delta N_{\rm eff}$ 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	+ `sn.pantheonplus`	frozen (Table II)
Full-tension	+ `H0.riess2020Mb` + DES-Y3 S_8 Gaussian	frozen (Table II)

VI. COSMIC BIREFRINGENCE: SPECTATOR ALP CONSISTENCY CHECK

Framing (prior-volume accommodation, not an independent test).—**The accommodation fraction reported here is a prior-predictive / prior-sensitivity exercise at fixed $C_{a\gamma}$, not a measurement and not a distinctive ECH prediction.** This section is an *accommodation / prior-volume* exercise, not an independent confirmation of the birefringence signal. The ALP likelihood is a Gaussian summary centred on the single published datum $\beta_{\rm obs} = 0.342^\circ \pm 0.094^\circ$ [13], so the posterior β values ($\beta_{\rm ALP}$, $\beta_{\rm free}$) agreeing with $\beta_{\rm obs}$ "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 weighted median $m = 37.2\, H_0$ well above the scan-prior midpoint — and (ii) the location of the controlled spectator-validity region: 13.3818% of the likelihood-weighted samples in this fixed-background surrogate chain satisfy $\Omega_a < 0.01$. This is a chain-selection frequency, not a physical posterior probability or a prior-volume cost. Within this controlled subset, the conditional median is $\theta_i = 0.211$ (16–84%: 0.149–0.267). The much smaller $\theta_i \leq 0.1$ sliver is a separate, low-sample diagnostic, not a description of the $\Omega_a < 0.01$ subset. 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.

Prior-predictive method.—The reported 11.6% and 6.1% fractions come from N = 100,000 independent draws per configuration with NumPy seed 1234 and no likelihood weighting (`reproducibility/cosmology/alp_prior_predictive.py`). In both configurations $\theta_i \sim \mathcal{U}[0.01, \pi]$, $\log_{10}(m_a/\text{eV}) \sim \mathcal{U}[-35, -30]$, and $f_a = M_{\rm Pl}$; the fixed-coupling run sets $C_{a\gamma} = 8$, while the broad-coupling run draws $C_{a\gamma} \sim \mathcal{U}[4, 60]$. Each draw is propagated through the committed nonlinear equation of motion using an adaptive DOP853 estimator whose maximum discrepancy from the reference integrator over the validation sample is 2.49×10^{-8} degrees. Counting uses the signed band $|\beta - 0.342^\circ| < 0.094^\circ$ (not an unsigned-tail statistic): 11.597% of fixed-coupling draws and 6.137% of broad-coupling draws fall inside it. Their binomial Monte Carlo standard errors are respectively 0.101 and 0.076 percentage points. The complete result, including the 2σ counts, is archived at `reproducibility/cosmology/alp_prior_predictive_result.json`.

Applying the ODE-derived spectator criterion $\Omega_a < 0.01$ to the same seed-1234 trajectories retains 34,947/100,000 (34.947%; MC standard error 0.151 percentage points). Within this retained subset, the fixed-coupling arm has zero observed draws (0/34,947) in the signed 1σ band (95% Wilson upper bound 0.011%), while the broad-coupling arm has 863/34,947 = 2.469% (MC standard error 0.083 percentage points). The unrestricted values above are therefore fixed-background ALP prior-predictive fractions, whereas these latter values are explicitly conditioned on spectator validity. Neither is Bayesian evidence, a posterior probability, or a physical prior cost; the conditioned calculation remains a prior-frequency diagnostic under the same fixed-background equations. Its deterministic generator, result, and execution receipt are archived at `reproducibility/cosmology/alp_spectator_conditioned_prior_predictive.py`, `reproducibility/cosmology/alp_spectator_conditioned_prior_predictive_result.json`, and `reproducibility/cosmology/alp_spectator_conditioned_prior_predictive_receipt.json`.

Note (spectator-status caveat, main text).—Spectator validity is defined by the derived energy-density cut $\Omega_a < 0.01$, not by an imposed small-angle threshold. The conditional median $\theta_i = 0.211$ is distinct from the sparsely sampled $\theta_i \leq 0.1$ sliver. The often-quoted factor $[0.5/0.1]^2 = 25$ is only an energy-density scaling ratio at fixed mass, not a tuning probability or prior-volume cost. The operational criterion for spectator status throughout this section is $\Omega_a < 0.01$: 13.3818% of the likelihood-weighted samples in the fixed-ΛCDM-background surrogate chain satisfy this cut (Table V). This fraction is not a physical posterior probability because samples with non-negligible Ω_a were nevertheless evaluated on an unchanged background. The conditional $\Omega_a < 0.01$ readout is the controlled subset for which that approximation is self-consistent at the percent level. The threshold value 0.01 is physically motivated rather than tuned to a target: $\Omega_a = \rho_a(z{=}0)/\rho_{\rm crit,0}$ 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 V) lists both the $\Omega_a < 0.1$ (44.0474% surrogate-chain weight, $\mathrm{ESS}_\beta = 1597$) and $\Omega_a < 0.01$ (13.3818%, $\mathrm{ESS}_\beta = 485$) readouts. The rotation marginal remains compatible within its uncertainty across them (the verdict does not flip between the two thresholds); we do not attempt a Bayes-factor cut, which would require a dedicated model-comparison analysis outside the scope of this consistency check. This subsection presents a secondary consistency check using the spectator ALP model. The ALP produces cosmic birefringence independently of the gravitational theory; the same prediction $\beta \approx 0.27^\circ$ arises in standard GR with an identical ALP Lagrangian within the scan-prior envelope but near its upper-displacement/coupling edge; the posterior-supported fixed-$C_{a\gamma} = 8$ fit shifts to $m \gg H_0$ (weighted median $37.2\,H_0$). 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.* [22].

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) [13] (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 [12], quoted as an independent literature cross-check), which accounts for shared calibration systematics. We do not form a naive cross-dataset significance from these differently constructed summaries.

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

$$\Delta\phi/f_a \approx 0.42 \qquad (m = 2H_0,\ \theta_i = 1). \tag{8}$$

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 $z_{\rm init} = 3000$ reproduces $\Delta\phi/f_a = 1.0601$ to four significant figures):

$$\begin{aligned} \beta &\approx \frac{\alpha_{\rm EM}}{4\pi}\, C_{a\gamma}\, \frac{\Delta\phi}{f_a} \\ &= \underbrace{(7.297 \times 10^{-3}/(4\pi))}_{\alpha_{\rm EM}/(4\pi)\,\approx\, 5.81\times 10^{-4}} \times 8\, \times\, 1.06 \\ &= 4.93 \times 10^{-3}\,\mathrm{rad} \times \frac{180^\circ}{\pi} \ \approx\ 0.28^\circ. \end{aligned} \tag{9}$$

The product of the three factors is $5.81 \times 10^{-4} \times 8 \times 1.06 = 4.93 \times 10^{-3}\,\mathrm{rad}$ (with $\alpha_{\rm EM} = 1/137.036$; the four-figure committed Runge–Kutta value $\Delta\phi/f_a = 1.0601$ changes only the displayed higher precision, not this rounded result). The $\alpha_{\rm EM}/(4\pi)$ prefactor is convention-dependent ($\alpha_{\rm EM}/(2\pi)$ 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 $g_{a\gamma} = C_{a\gamma}\,\alpha_{\rm EM}/(2\pi f_a)$ and $\beta = (g_{a\gamma}/2)\,\Delta\phi$, the convention fixed in the committed model specification and implemented identically throughout the numerical pipeline (`research/branch_R_alp_birefringence/02_birefringence_prediction.md`; `research/branch_R_alp_birefringence/phase2_mcmc/alp_ode.py`). The fiducial value $\beta \approx 0.27^\circ$ corresponds at $C_{a\gamma} = 8$ to EOM trajectories with, e.g., ($\theta_i = 1$, $m \approx 3.9\,H_0$) or ($\theta_i \approx 1.4$, $m \approx 3\,H_0$). The committed EOM integration (`research/branch_R_alp_birefringence/phase2_mcmc/alp_ode.py`) gives $\Delta\phi/f_a = 0.35$–0.42 at masses $m \approx 1.8$–$2\,H_0$ ($\theta_i = 1$); 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 $\beta \approx 0.01$–0.48° (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 $\theta_i = 2$,

---

[5] The ALP ODE is integrated on a ΛCDM 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 [23]. The ΛCDM-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.

[6] At fixed mass and in the small-angle limit, Ω_a scales approximately as θ_i^2, so changing θ_i from 0.5 to 0.1 lowers the energy-density estimate by a factor 25. This is an energy-density ratio, not a tuning probability. The executed chain classifies spectator status directly through the nonlinear-ODE value $\Omega_a < 0.01$; that subset has median $\theta_i = 0.211$, while $\theta_i \le 0.1$ is a separate indicative sliver containing only 42 raw samples.

$m = 3H_0$; 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. Spectator status is assigned from the ODE-derived $\Omega_a < 0.01$ cut (conditional median $\theta_i = 0.211$), not from an imposed $\theta_i = 0.1$ threshold. The model *accommodates* the observed signal; the reported selection frequency is not a tuning probability. The $[0.01, 0.48]°$ 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 $\beta_{\rm ALP} \in [0.01, 0.48]°$ across the full scan-prior box (the benchmark EOM grid, not the spectator-consistent subset).

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 $\beta_{\rm free}$, 720; chain configurations and per-chain counts in Appendix C) yields: $\beta_{\rm ALP} = 0.336° \pm 0.10°$ ($C_{a\gamma} = 8$ 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, $\beta \equiv \beta + n \times 90°$ for E/B, and the Gaussian summary ignores wrapping — harmless here because the posterior support is confined to $|\beta| \lesssim 0.7° \ll 90°$, so wrapped images carry negligible likelihood), has weighted median and 16–84% interval $\theta_i = 1.317^{+0.439}_{-0.369}$ and $m \sim 10$–$10^2\,H_0$ (weighted posterior median $m = 37.2\,H_0$) — 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 $\Delta\phi/f_a \approx 1.29$, ∼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 $\beta_{\rm obs} = 0.342°$), consistent with the model-independent fit $\beta_{\rm free} = 0.344° \pm 0.10°$ (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 $\beta_{\rm obs} = 0.342° \pm 0.094°$ [13] — not a re-analysis of the *EB* spectra themselves; 720 accepted samples in the dedicated $\beta_{\rm free}$ configuration; full priors and likelihood details in Appendix C; $\beta_{\rm free}$ denotes the unconstrained-amplitude fit distinct from $\beta_{\rm ALP}$ which has $C_{a\gamma} = 8$ fixed) and the observed $\beta_{\rm obs} = 0.342° \pm 0.094°$. All three within 1σ (e.g. $\beta_{\rm ALP}$ vs $\beta_{\rm obs}$: the model posterior mean $0.336°$ lies 0.06σ from the data central value, $\Delta = 0.006°$, using the data uncertainty $\sigma_{\rm obs} = 0.094°$; these are consistency statements against the same single published measurement, not independent confirmations). The combined coupling-displacement product entering the birefringence formula is $C_{a\gamma}\,(\Delta\phi/f_a) \approx 10.3$ ($\beta = 0.342°$ in radians is 5.97×10^{-3}, the prefactor $\alpha_{\rm EM}/(4\pi)$ is 5.8×10^{-4}, giving $C_{a\gamma}\,\Delta\phi/f_a = \beta/[\alpha_{\rm EM}/(4\pi)] \approx 10.3$); 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 ($m \sim 10$–$10^2\,H_0$, where the displacement saturates near $\Delta\phi/f_a \approx 1.2$–$1.3$), the required coupling is ≈8–10. 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 the non-minimal photon-coupling burden is independent of the ODE-derived spectator cut: $C_{a\gamma} \gtrsim 9$ is required across the reported region. 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 ($\gtrsim$ 50–160) 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. At fixed $\beta = 0.342°$, $\Delta\phi/f_a \propto \theta_i$ along the underdamped trajectory, so a 5× 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 ∼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% $[6.8, 43.4]$; $\beta = 0.328° \pm 0.100°$) 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 $\beta = 0.328° \pm 0.100°$ posterior matches the observed $0.342° \pm 0.094°$, confirming the consistency-check verdict over the posterior-supported coupling range rather than at fixed benchmarks. *Spectator-subset readout* (same chain, no additional sampling): the chain's derived ALP energy fraction Ω_a (defined in the subsection below) is computed by integrating the nonlinear cosine-potential equation of motion to today and evaluating the final kinetic plus potential energy. These per-sample ODE values satisfies $\Omega_a < 0.1$ for 44.0474% and $\Omega_a < 0.01$ for 13.3818% of the fixed-background surrogate-chain weight; restricted to the $\Omega_a \leq 0.01$ spectator-safe subset, the rotation marginal is $\beta = 0.2759° \pm 0.0988°$ (weighted mean±standard deviation; weighted median 0.2841°). The injection fiducial 0.27° is the rounded value of the canonical injection point, distinct from this posterior marginal, which is consistent with $\beta_{\rm obs} = 0.342° \pm 0.094°$ at 0.5σ. *H_0 marginalization note:* Ω_a at each MCMC step uses the ΛCDM background Hubble rate evaluated at the Cobaya posterior mean $H_0 = 67.68\,{\rm km\,s^{-1}\,Mpc^{-1}}$ (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 0.3275% of the surrogate-chain 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 describes a separate extreme small-angle sliver; it is not the $\Omega_a < 0.01$ subset, whose conditional median is $\theta_i = 0.211$. Convergence: $\hat{R} - 1 < 0.01$ for all runs. The joint posterior structure of this continuous-prior configuration is shown in Fig. 3. 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 $\Omega_a = \rho_a(z=0)/\rho_{\rm crit,0}$. We define this quantity precisely here because the Table V entries (44.0474% at $\Omega_a < 0.1$; 13.3818% 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

$$V(\phi) = m_a^2 f_a^2 \left[1 - \cos\left(\frac{\phi}{f_a}\right)\right], \tag{10}$$

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 $\mathcal{O}(\theta_i^2/12)$.

Nonlinear evolution and energy density today.—At each MCMC step, the implemented module receives (m_a, θ_i) with $f_a = M_{\rm Pl}$ fixed and integrates the full cosine-potential equation

$$\ddot{\theta} + 3H\dot{\theta} + m_a^2 \sin\theta = 0 \tag{11}$$

from $z = 3000$ to $z = 0$ on a fixed ΛCDM background. The solver uses $\ln a$ as the independent variable, DOP853 with relative tolerance 10^{-10}, absolute tolerance 10^{-12}, and initial conditions $\theta(z = 3000) = \theta_i$, $d\theta/d\ln a = 0$; see `research/branch_R_alp_birefringence/phase2_mcmc/alp_ode.py`. From the final state it evaluates

$$\Omega_a \equiv \frac{\rho_a(0)}{\rho_{\rm crit,0}}, \qquad \rho_a(0) = \frac{1}{2} f_a^2 \dot{\theta}_0^2 + m_a^2 f_a^2 (1 - \cos\theta_0), \tag{12}$$

with $\rho_{\rm crit,0} = 3H_0^2 M_{\rm Pl}^2$. Thus the chain's Ω_a column is an ODE-to-today kinetic-plus-potential result; no $3H = m_a$ onset rule or matter-dilution approximation is implemented in the load-bearing calculation. 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 V below. The 44.0474%/13.3818% figures are direct weighted selection frequencies over these ODE-derived per-step values. Because the likelihood calculation held the ΛCDM background fixed even when the derived Ω_a was appreciable, they are neither physical posterior probabilities nor prior costs. They are conditional on the stated priors and surrogate likelihood. Only the $\Omega_a < 0.01$ conditional subset is used for a controlled spectator readout.

LiteBIRD forecast.—LiteBIRD is projected to achieve $\sigma(\beta) \approx 0.03°$ [24] (a forecast under foreground- and

TABLE V. Fixed-background surrogate-chain readout of the continuous-prior `c5_continuous` chain on three spectator-status subsets, together with the full chain. The selection-fraction column gives the fraction of likelihood-weighted surrogate-chain MC weight (and, for the strict $\theta_i \leq 0.1$ sliver, raw-count fraction in parentheses). β is the weighted mean±standard deviation; m/H_0, θ_i, and $C_{a\gamma}$ are weighted 16/50/84 percentiles. Every ESS uses the same marker and estimator: $\beta_{\rm deg}$ after integer-weight expansion, per-chain FFT autocorrelation, Sokal automatic window $c = 5$, and summed per-chain N/τ. The exact generator, input hashes, per-chain diagnostics, and compact output are in `c5_table_iv_recompute.py/json`. The $\Omega_a < 0.01$ subset is controlled at the percent-background level; the $\theta_i \leq 0.1$ row is indicative because it contains only 42 raw samples. Selection fractions are not physical posterior probabilities: they are conditional on the fixed-background surrogate and Gaussian summary likelihood for the Eskilt–Komatsu $\beta_{\rm obs}$ datum (Appendix C) and would shift under a self-consistent background and full-EB re-analysis.

Subset	sel. fraction	β (deg)	m/H_0	ESS_β
full chain	100%	0.3263 ± 0.0990	1.71/13.17/196.65	2949
$\Omega_a < 0.1$	44.0474%	0.3147 ± 0.1026	4.70/37.78/264.48	1597
$\Omega_a < 0.01$ (controlled)	13.3818%	0.2759 ± 0.0988	5.99/40.74/239.01	485
$\theta_i \leq 0.1$ (indicative)	0.3275% (0.4690% raw)	0.1172 ± 0.0466	4.71/41.53/194.14	95

Subset	θ_i (16/50/84)	$C_{a\gamma}$ (16/50/84)
full chain	0.32/0.97/2.16	7.34/20.74/45.63
$\Omega_a < 0.1$	0.22/0.41/0.70	14.23/26.21/46.41
$\Omega_a < 0.01$ (controlled)	0.15/0.21/0.27	29.89/43.31/54.14
$\theta_i \leq 0.1$ (indicative)	0.061/0.083/0.096	36.82/47.18/55.64

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

Caveats.—This birefringence fit is independent of bounce cosmology: the ALP is a spectator field that does not participate in bounce dynamics. The benchmark $f_a \sim M_{\rm Pl}$ is a prior choice; no derived photon-torsion coupling connects the Holst action to this 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 $m \gg H_0$ (weighted median $37.2\, H_0$). The spectator-safe interpretation is conditional on the ODE-derived $\Omega_a < 0.01$ subset and requires a non-minimal photon coupling; it does not uniquely predict the signal.

VII. CONCLUSIONS

This paper documents three standalone technical reproducibility and consistency-check analyses, with Paper I(a) cited only for the separate minimal-action theory context [1]. To restate the scope plainly before the per-analysis summaries: none of the three *verifies* the ECH spin-torsion sector itself. The $\Delta N_{\rm eff}$ 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 characterization, 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."

ΛCDM+$\Delta N_{\rm eff}$ 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 $\Delta N_{\rm eff}$ consistent with zero in both frozen dataset combinations and recovers $H_0 = 67.68 \pm 1.06\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$, consistent with standard Planck ΛCDM. The $\Delta N_{\rm eff}$ extension does not resolve the Hubble tension. We note that the $\Delta N_{\rm eff}$ 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 sensitivity rerun; its non-converged 0.04σ central-value agreement is corroborative only and does not bound pairing-induced bias (see Sec. V, *Release-pairing note*). Current data neither require nor exclude a phenomenological $\Delta N_{\rm eff}$ extra-radiation contribution; CMB-S4 ($\sigma(N_{\rm eff}) \sim 0.03$ [10]) 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 ($\Delta N_{\rm eff}$, H_0, etc.) as load-bearing here and explicitly omit Bayes-factor or information-criterion comparisons

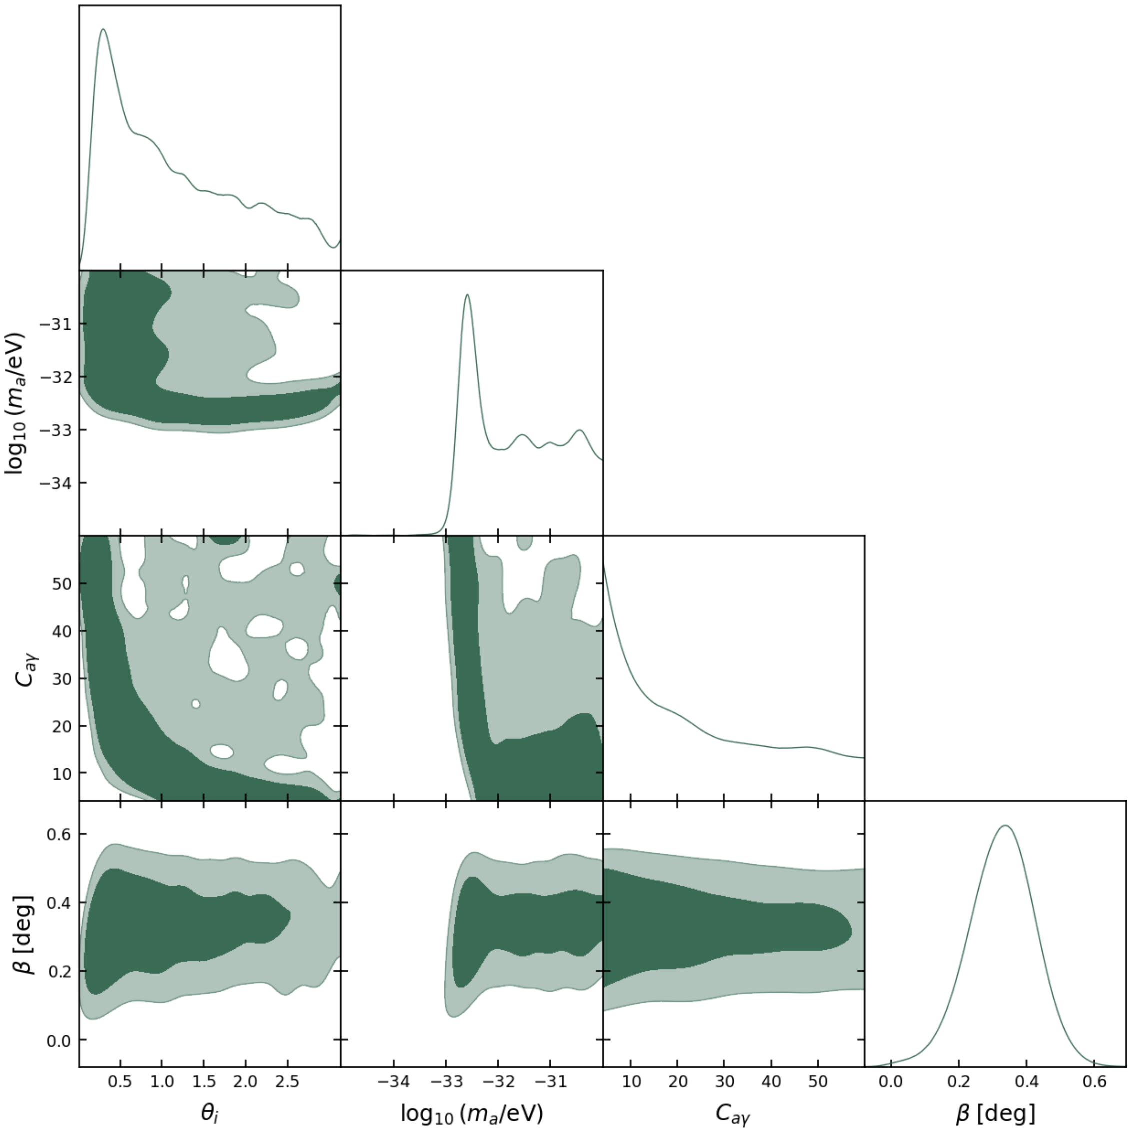

FIG. 3. Spectator-ALP joint posterior triangle from the continuous-prior cross-check configuration in which the photon anomaly coefficient is sampled freely (flat priors $C_{a\gamma} \in [4, 60]$ — shifted and extended from the earlier $[1, 30]$ to cover the posterior-supported coupling band (median 20.7, 16–84% $[7.3, 45.6]$); the dropped $[1, 4)$ interval lies entirely below the minimum coupling $\approx$8.6 required to reach the *central* value $\beta = 0.342°$ 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 — $\theta_i \in [0.01, \pi]$, $\log_{10}(m_a/\mathrm{eV}) \in [-35, -30]$ — the mass prior corresponds to $m/H_0 \approx 7\times10^{-3}$ to 7×10^{2} for $H_0 = 67.7$ km/s/Mpc $= 1.44\times10^{-33}$ 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, $\beta = 0.326° \pm 0.099°$, is consistent within 1σ with the fixed-$C_{a\gamma} = 8$ result $\beta_{\mathrm{ALP}} = 0.336° \pm 0.10°$ and the observed $\beta_{\mathrm{obs}} = 0.342° \pm 0.094°$ (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).

in this Metropolis-Hastings-only analysis.

NaMaster pipeline validation.—The 500-MC pseudo-C_ℓ analysis on synthetic ΛCDM polarization skies confirms that the deconvolution pipeline recovers injected birefringence angles with a canonical MC-mean bias $\hat\beta - \beta_{\mathrm{inj}} = -0.0010° \pm 0.0023°$ (standard error; Sec. IV). The diagonal-covariance template score is a heuristic that omits inter-bin covariance. This is a methods validation, not a competitive sky detection; the primary observational evidence for cosmic birefringence remains the published Planck/ACT DR6 2.7–2.9σ measurements.

Spectator-ALP consistency.—To state the verdict first: the spectator ALP *accommodates*, it does not *predict*, the observed birefringence, and the controlled spectator readout rests on the $\Omega_a < 0.01$ subset (13.3818% of the fixed-background surrogate-chain weight), not on the extreme $\theta_i \leq 0.1$ sliver (0.3275%, indicative only). These are selection frequencies, not physical posterior probabilities. An independent unweighted prior-predictive calculation gives a different, prior-space question: 34,947/100,000 trajectories satisfy $\Omega_a < 0.01$, of which zero are observed at fixed $C_{a\gamma} = 8$ (0/34,947; 95% Wilson upper bound 0.011%) and 863 for $C_{a\gamma} \sim \mathcal{U}[4, 60]$ (2.469%; MC standard error 0.083 percentage points) also lie in the published signed 1σ band. These conditional Monte Carlo frequencies are neither evidence nor posterior probabilities and are distinct from the likelihood-weighted surrogate-chain selection frequency above. Within that $\Omega_a < 0.01$ spectator-safe subset (13.3818% surrogate-chain weight; median $m = 40.74\,H_0$, $\Omega_a < 0.01$, $\mathrm{ESS}_\beta = 485$), an ALP with $f_a \sim M_{\mathrm{Pl}}$ and conditional median $m = 40.74\,H_0$ (not $m \sim H_0$; the scan-prior envelope spans $m \sim H_0$ but the conditional fit shifts well above the prior midpoint) is consistent with the published 3.6σ joint signal, with the caveat that this spectator-safe subset has conditional median $\theta_i = 0.211$ (16–84%: 0.149–0.267); the separate $\theta_i \leq 0.1$ sliver is indicative only (fn. 6). It satisfies spectator status ($C_{a\gamma}\,\Delta\phi/f_a \approx 10.3$ for $\beta = 0.342°$, 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; no ECH interpretation is required. Under forecast foreground-cleaning and calibration assumptions [24], a $\beta \simeq 0.27°$ signal would be detected relative to $\beta = 0$ (null-rejection) at $\sim 9\sigma$ by LiteBIRD in the early 2030s, while the model-discrimination separation from the published central value is only ≈0.7σ (distinct null hypotheses; see Sec. VI).

Data and Code Availability

All materials are at:

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

This pre-release revision is identified by the in-tex `v1B.0.109` stamp (defined in `paper1b_mcmc_companion.tex` `\newcommand{\paperVersion}`). It does *not* claim that an earlier repository commit freezes the new paper. Exact analysis-artifact provenance is instead recorded in `reproducibility/p1b_analysis_artifact_manifest_v1B.0.108.json`. That manifest records the base repository commit only as context and explicitly marks itself pre-release; the eventual submission commit/tag will freeze the paper source, manifest, and compiled PDF together. For ordinary Git blobs the manifest records the SHA-256 and byte size of the scientific file present in the repository. For frozen-chain `spin_torsion.1.txt` entries, however, the local 133-byte Git blob is a Git LFS pointer: its local SHA-256 authenticates only that pointer record. The manifest therefore records separately the pointer SHA-256 and the scientific payload's Git LFS `oid sha256:...` and declared byte size. The payload bytes are served by the MCMC-diagnostics HuggingFace mirror listed in Appendix A; payload verification means matching that LFS OID and size after retrieval, not equating the pointer hash with the chain. 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 [25] but not by any analysis of this paper. Frozen MCMC chain pointer records, payload OIDs, and declared payload sizes are committed; the scientific payloads are mirrored on HuggingFace (Appendix A). Fresh re-verification chains can be regenerated via `reproduce_cosmology.sh` (∼4–12 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.0.109`; 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. *AI-assisted methods.*—AI-assisted coding and language tools were

used, under the author's direction, for MCMC configuration checks, pipeline scripting, symbolic and dimensional cross-checking, and manuscript preparation. Vendor and model-version names are omitted because they are not scientific dependencies; the executable environments and load-bearing results are fixed by the committed code, data, and SHA-256 manifest. All scientific claims, derivations, and numerical results were verified by the author against the committed computational artifacts—Cobaya configurations, frozen chains, the NaMaster recovery suite, and the ALP prior-predictive outputs—documented in the public reproducibility manifest (Appendix A); together with the repository git history these constitute a public audit trail. The author takes sole responsibility for all scientific claims, derivations, numerical results, and bibliographic attributions. 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 [26] 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.
- `reproducibility/cosmology/alp_spectator_conditioned_prior_predictive.py` and its adjacent result and execution-receipt JSON files—the deterministic 100,000-draw-per-arm unrestricted and $\Omega_a < 0.01$-conditioned ALP prior-predictive calculation.

What is included vs. regenerable.—The two frozen ΛCDM+$\Delta N_{\rm eff}$ directories are committed in-repo. Their small diagnostics are ordinary Git payloads; each large chain file is represented locally by a Git LFS pointer whose scientific payload OID and byte size are recorded in the v1B.0.108 manifest, with the scientific bytes mirrored in the MCMC-diagnostics HuggingFace dataset. Together these artifacts back Table II, Table IV, and the relevant Table VI claim-classification entries of this paper; the ALP $c5$ continuous chain backs Table V 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 ΛCDM+$\Delta N_{\rm eff}$ proxy chains for independent re-verification are NOT bundled and must be regenerated locally via `reproduce_cosmology.sh` ($\sim$4–12 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; no DOI is fabricated here. In the interim, the stable references are the frozen-chain Git LFS payload OIDs/sizes, HuggingFace payload mirrors, and ordinary-file repository SHAs pinned in `CHANGELOG.md` and `reproducibility/`; these are version-locked without fabricating a DOI. The URLs are also recorded in the repository `CHANGELOG.md` under the entry for `v1B.0.109` and are preserved against future README changes via the version-stamp commit identified by the changelog):

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

Appendix B: Claims Classification

Table VI 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.

TABLE VI. 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 ordinary Git artifacts or from scientific Git LFS payloads identified by payload OID and size in the manifest for `v1B.0.109`; a paper-freezing public commit/tag is pending (see Data and Code Availability).

Claim	Type	Reference value	Notes
$\Delta N_{\rm eff} = -0.020 \pm 0.169$ (full-tension)	MCMC	Int. verified (frozen chains)	Stock CAMB proxy
$\Delta N_{\rm eff} = +0.058 \pm 0.179$ (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 ΛCDM
$H_0 = 67.78 \pm 1.09$ (Planck+BAO+SN)	MCMC	Int. verified; Planck-consistent	Recovers ΛCDM
Model-comparison ΔAIC/BIC/$\ln B$	Numerical	Not reported	Nested-sampling follow-up
$\hat{\beta}_{\rm NaMaster} = 0.269°$ (500-MC)	Numerical	Int. verified (pipeline)	Exact-window bias consistent with zero
$\beta_{\rm ALP} = 0.336° \pm 0.10°$	MCMC	Int. verified (ALP chains)	ALP MCMC
Published 3.6σ ($\beta = 0.342 \pm 0.094°$)	Lit.	[13]	Eskilt et al.
Stock CAMB proxy $\neq$ ECH theory module	Scope	Definition	§III
ALP birefringence not distinctive ECH prediction	Scope	Definition	§VI

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

The ALP-MCMC results quoted in Sec. VI ($\beta_{\rm ALP} = 0.336° \pm 0.10°$ at $C_{a\gamma} = 8$ fixed; $\beta_{\rm free} = 0.344° \pm 0.10°$ 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]
- $\log_{10}(m_a/\mathrm{eV})$: uniform prior on $[-35, -30]$ ($m/H_0 \approx 7 \times 10^{-3}$ to 7×10^2).
- f_a: fixed at $M_{\rm Pl}$ (spectator-class theoretical input from the Holst-sector pseudoscalar structure; not sampled).

This configuration yields $\beta_{\rm ALP} = 0.336° \pm 0.10°$ with weighted posterior median and 16–84% interval $\theta_i = 1.317^{+0.439}_{-0.369}$ and weighted median $m = 37.2\,H_0$.

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 ∼28% of the coupling posterior mass.

Configuration (iii) — model-independent $\beta_{\rm free}$ fit (`run3_baseline`; 720 accepted samples).

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

Sampled parameters and priors (continuous-prior cross-check configuration; Fig. 3).—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 $[0.01, \pi]$.
- $\log_{10}(m_a/\mathrm{eV})$: 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 applied directly to the committed chain columns):

[7] Backreaction / spectator-status disclosure: see fn. 6 in Sec. VI. The wide prior is an envelope-completeness choice. Spectator status is classified from the ODE-derived $\Omega_a < 0.01$ cut, whose conditional median is $\theta_i = 0.211$; the separate $\theta_i \leq 0.1$ row is indicative only.

Parameter	Chain	$N_{\rm acc}$	ESS
θ_i	`c5_continuous`	8,955	796
$\log_{10}(m_a/\mathrm{eV})$	`c5_continuous`	8,955	1,114
$C_{a\gamma}$	`c5_continuous`	8,955	814
$\beta_{\rm deg}$ (derived)	`c5_continuous`	8,955	2,949
$\beta_{\rm free}$	`run3_baseline`	720	265

The `c5_continuous` ESS values (∼800–2949) are adequate for posterior characterisation. The $\beta_{\rm free}$ run3_baseline chain (720 accepted samples, ESS≈ 265) is marginal; the result $\beta_{\rm free} = 0.344^\circ \pm 0.10^\circ$ 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 $\beta_{\rm obs} = 0.342^\circ \pm 0.094^\circ$ [13] (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 $(\beta_{\rm obs}, \sigma_\beta)$ 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 could change the location, width, and tails of the amplitude posterior: the unrestricted fixed-background fraction (11.6% within 1σ at fixed $C_{a\gamma} = 8$), its spectator-conditioned counterpart (zero observed, 0/34,947; 95% Wilson upper bound 0.011%), and the $\Omega_a < 0.01$ fixed-background surrogate-chain selection frequency therefore inherit this approximation. The latter is not a physical posterior probability or prior cost. The reported posterior medians ($m = 37.2\,H_0$, $C_{a\gamma}\,\Delta\phi/f_a \approx 10.3$) are conditioned on that summary likelihood. A full joint-EB re-fit is required to determine the direction and size of any shifts; no robustness forecast for those shifts is made here. The separate ACT DR6 measurement $\beta = 0.215^\circ \pm 0.074^\circ$ [12] 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) ($N_{\rm tot} = 9{,}720$ 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*: $\beta_{\rm ALP} = 0.336^\circ \pm 0.10^\circ$ — the fixed-$C_{a\gamma} = 8$ chain `run1_full` (2,160 samples); $\beta_{\rm free} = 0.344^\circ \pm 0.10^\circ$ — 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 ($m \sim H_0$, $f_a \sim M_{\rm Pl}$; the $C_{a\gamma} \in [4, 12]$ benchmark sweep plus the continuous-prior $C_{a\gamma} \in [4, 60]$ headline configuration). The consistency with $\beta_{\rm obs}$ 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, Algebraic Torsion in Minimal Einstein–Cartan–Holst Gravity: Four-Fermion Contact Benchmarks and Classical Scalar-Sector Transparency (2026), companion paper, posted concurrently on arXiv [arXiv:XXXX.XXXXX — ID inserted at coordinated submission].

[2] 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].

[3] 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).

[4] N. J. Popławski, Four-fermion interaction from torsion as dark energy, General Relativity and Gravitation **44**, 491 (2012), arXiv:1102.5667 [gr-qc].

[5] 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].

[6] 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].

[7] 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].

[8] J. I. Kapusta and C. Gale, *Finite-Temperature Field Theory: Principles and Applications*, 2nd ed., Cambridge Monographs on Mathematical Physics (Cambridge University Press, 2006).

[9] M. Brüggen, Effects of a torsion field on Big Bang nucleosynthesis, General Relativity and Gravitation **31**, 1935 (1999), arXiv:astro-ph/9906403 [astro-ph].

[10] CMB-S4 Collaboration, K. Abazajian, *et al.*, CMB-S4 Science Book, First Edition, arXiv e-prints (2016), arXiv:1610.02743 [astro-ph.CO].

[11] P. Diego-Palazuelos, J. R. Eskilt, Y. Minami, M. Tristram, *et al.*, Cosmic birefringence from the Planck data release 4, Phys. Rev. Lett. **128**, 091302 (2022), reports beta = 0.30 +/- 0.11 deg from Planck NPIPE (PR4), arXiv:2201.07682 [astro-ph.CO].

[12] 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].

[13] 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].

[14] 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].

[15] Planck Collaboration, N. Aghanim, *et al.*, Planck 2018 results. VI. cosmological parameters, Astronomy & Astrophysics **641**, A6 (2020), arXiv:1807.06209 [astro-ph.CO].

[16] 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].

[17] 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].

[18] 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].

[19] D. Brout *et al.*, The Pantheon+ analysis: Cosmological constraints, Astrophys. J. **938**, 110 (2022), arXiv:2202.04077 [astro-ph.CO].

[20] 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].

[21] 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].

[22] 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].

[23] 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].

[24] 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].

[25] H. Golden, Galaxy Chirality at Scale: 8.47M Galaxies Classified, Hemisphere Null at $p_{\rm LEE} < 10^{-4}$ (2026), companion paper, posted concurrently on arXiv.

[26] 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].