

Testing parity symmetry with the polarized cosmic microwave background

Oliver H. E. Philcox^{1,2,*} and Maresuke Shiraishi^{3,†}

¹*Center for Theoretical Physics, Department of Physics, Columbia University, New York, New York 10027, USA*

²*Simons Society of Fellows, Simons Foundation, New York, New York 10010, USA*

³*School of General and Management Studies, Suwa University of Science, Chino, Nagano 391-0292, Japan*



(Received 14 August 2023; accepted 18 March 2024; published 11 April 2024)

New physics in the early Universe could lead to *parity-violation* in the late Universe, sourcing statistics whose sign changes under point reflection. The best constraints on such phenomena have come from the *Planck* temperature fluctuations; however, this is already cosmic-variance limited down to relatively small scales. Thus only small improvements are expected in the future. Here, we search for signatures of parity violation in the polarized cosmic microwave background, using the *Planck* PR4 *T*- and *E*-mode data. We perform both a simulation-based blind test for any parity-violating signal at $\ell < 518$, and a targeted search for primordial *U*(1) gauge fields (and the amplitudes of a generic collapsed model) at $\ell < 2000$. In all cases, we find no evidence for new physics, with the model-independent test finding consistency with the FFP10/NPIPE simulation suite at $(-)$ 0.4 σ , and the gauge field test constraining the fractional amplitude of gauge fields during inflation to be below 6×10^{-19} at 95% confidence level for a fiducial model. The addition of polarization data can significantly improve the constraints, depending on the particular model of primordial physics, and the bounds will tighten significantly with the inclusion of smaller-scale information.

DOI: [10.1103/PhysRevD.109.083514](https://doi.org/10.1103/PhysRevD.109.083514)

I. INTRODUCTION AND MOTIVATION

By measuring the large-scale temperature anisotropies present in the cosmic microwave background (CMB), satellites such as *Planck* and WMAP have transformed our view of the primordial Universe [e.g., [1–5]]. To some extent, this mission is already complete: due to the cosmic variance limit, the future will not yield additional information on the Universe from large scales ($\ell \lesssim 2000$). Polarization, on the other hand, provides new clues to terrestrial cosmic detectives. In *Planck*, the cosmic variance limit impacts *E*-modes only on the largest scales, and investigations into *B*-modes are hampered only by noise and lensing effects. As such, precision analyses of polarization will take the forefront in the search for fundamental physics from the CMB; indeed, this is already occurring with observatories such as ACT, SPT, and BICEP/Keck [6–8].

To date, a wealth of new physics has been searched for in CMB temperature and polarization fluctuations, probing both the scalar and tensor sectors in the early and late Universe (see [4,5] and citations therein). A particularly novel example is that of parity violation; physics that

changes sign under point reflection. Although there exists a range of physical mechanisms to generate such effects [9–26], their signatures are difficult to probe in nature; if sourced by scalars, they appear only in four-point correlation functions and beyond ([12], see also [14,19,27,28]).¹ Difficult does not imply impossible however, and recent works have reported tentative signatures of parity violation in the late-time distribution of galaxies ([32,33], using the approach described in [27]) and in studies of CMB birefringence [34–36], though the latter requires late-time vector sourcing, and will not be considered in this work.

If the purported signal in the galaxy trispectrum were physical and primordial, what would it imply for CMB observations?² As discussed in [12,37], parity violation in the scalar sector imprints itself in projected two-sphere observables via a nonvanishing *T*- or *E*-mode trispectrum with odd $\ell_1 + \ell_2 + \ell_3 + \ell_4$. Geometrically, this can be thought of as a difference between trispectra of points with left- and right-handed chiralities. In the flat-sky

¹Physics sourced by gravitational waves, on the other hand, can be observed in lower-order statistics of *T*- and *E*-modes, e.g. bispectra with odd $\ell_1 + \ell_2 + \ell_3$ [4,29–31].

²See [21] for a discussion of potential late-time sources and their (in)viability.

*ohep2@cantab.ac.uk

†shiraishi_maresuke@rs.sus.ac.jp

and instantaneous-recombination limit, such effects vanish since observations are confined to a plane, and reflections become equivalent to rotations; in practice, this effect does not limit us, since it is relevant only when *all* sides of the tetrahedron (including its diagonals) are small (i.e. at high ℓ); a regime usually orthogonal to the signals of interest.

In [37], the trispectrum of the CMB temperature fluctuations was used to search for scalar parity violation. Despite the analysis containing $\approx 250\times$ more modes than the large-scale structure equivalent, no signal was found in a blind analysis, casting doubt on any primordial explanation of the former results. Here, we extend the analysis to the polarization sector [and additionally use updated *Planck* fourth data release (PR4) data [38–40]], making use of the new trispectrum estimators described in [41] (building on [42]). Noting that *E*-modes also trace primordial scalar physics, this can lead to a significant boost in signal-to-noise (depending on the model in question), particularly due to the large number of *T*- and *E*-mode trispectra that can be wrought (2^4 combinations) (cf. [4] for polarized bispectra). Our purpose is not to cast further aspersions on the large-scale structure model; rather, we wish to place novel bounds on little-known physics with high-precision data, which can be greatly improved with upcoming higher-resolution polarization experiments. We will further show the utility of performing targeted searches for models of interest, by constraining the gauge-field model described in [12]. This peaks in collapsed configurations, and its amplitude was seen at a $\approx 2\sigma$ hint in the temperature data [37].

The remainder of this paper is structured as follows. In Sec. II we discuss the *Planck* dataset, before describing the details of our trispectrum estimation pipeline in Sec. III. Sections IV and V present the main results of our analysis, general and targeted searches for primordial non-Gaussianity, before we conclude by discussing their implications in Sec. VI. Appendix A outlines the derivation of the gauge-field model used in Sec. V, whilst Appendix B presents various consistency checks of our estimators.

A. Conventions

Throughout the work, we define the (connected) trispectrum measured from a field $a_{\ell m}^X$ as

$$\left\langle \prod_{i=1}^4 a_{\ell_i m_i}^{X_i} \right\rangle_c = \sum_{LM} (-1)^M w_{\ell_1 \ell_2 m_1 m_2}^{L(-M)} w_{\ell_3 \ell_4 m_3 m_4}^{LM} f_{\ell_1 \ell_2, \ell_3 \ell_4}^{X_1 X_2, X_3 X_4}(L) + 23 \text{ perms}, \quad (1)$$

where $X_i \in \{T, E, B\}$, and L, M parametrize the internal leg of the tetrahedron [41,43]. This defines the weighting functions

$$w_{\ell_1 \ell_2 m_1 m_2}^{LM} = \sqrt{\frac{(2\ell_1 + 1)(2\ell_2 + 1)(2L + 1)}{4\pi}} \times \begin{pmatrix} \ell_1 & \ell_2 & L \\ -1 & -1 & 2 \end{pmatrix} \begin{pmatrix} \ell_1 & \ell_2 & L \\ m_1 & m_2 & M \end{pmatrix}, \quad (2)$$

which are akin to Gaunt integrals, but with modified spins. Trispectra containing an even number of *B*-modes are said to be *parity even* if $\ell_1 + \ell_2 + \ell_3 + \ell_4$ is even, and *parity odd* if $\ell_1 + \ell_2 + \ell_3 + \ell_4$ is odd; the situation is reversed if there is an odd number of *B*-modes. We assume a fiducial cosmology defined by the *Planck* 2018 parameters: $\omega_b = 0.022383$, $\omega_{\text{cdm}} = 0.12011$, $h = 0.6732$, $\tau = 0.0543$, $\log 10^{10} A_s = 3.0448$, $n_s = 0.96605$, for a single massive neutrino species with $m_\nu = 0.06$ eV [3]. To avoid confirmation bias, the pipeline discussed below was developed using only mock data, with the *Planck* results added as a final step.

II. DATA

Throughout this work, we utilize the PR4 from the *Planck* satellite [38–40], an upgrade from the *Planck* 2018 (PR3) data used in the previous work [37].³ This features improved treatment of polarization, noise, and large-scale systematics, which are particularly relevant to this study. In addition to the data, we use 600 end-to-end FFP10 simulations [38], with both data and mocks processed with the NPIPE pipeline, using the SEVEM component separation algorithm.⁴ These maps are smoothed to both $N_{\text{side}} = 256$ and $N_{\text{side}} = 1024$ resolution in HEALPix [44], facilitating different types of study.

When computing trispectra, we filter the data via a linear operator \mathbf{S}^{-1} . This firstly inpaints small holes and point sources in the map [containing less than $10(N_{\text{side}}/128)^2$ empty pixels] via a linear scheme [cf. [45]], before applying the *Planck* 2018 common component-separation mask [46], smoothed to $10'$ resolution, which retains $\approx 78\%$ of the sky.⁵ Secondly, we filter the data in harmonic space via a diagonal-in- ℓ weighting S_ℓ^{XY} , which is the sum of a theoretical (parity-even) power spectrum and a noise spectrum, extracted from the PR4 half-mission maps at $N_{\text{side}} = 2048$. We assume parity-even noise ($N^{EB} = 0$, $N^{EE} = N^{BB}$), and ignore any temperature-to-polarization leakage ($N^{TE} = N^{TB} = 0$). Finally, we include the beam (both in the filtering and the estimator); this matches the PR3 release,

³These have been made publicly available on NERSC, as described at portal.nersc.gov/project/cmb/planck2020/.

⁴This is preferred over the alternative COMMANDER algorithm, since the latter does not contain joint temperature and polarization simulations in PR4.

⁵Note that the mask here is included in the \mathbf{S}^{-1} weights, rather than the estimator itself [cf., [37,41]]. This leads to reduced leakage between bins and does not induce bias. Our approach matches [47] for the scalar case.

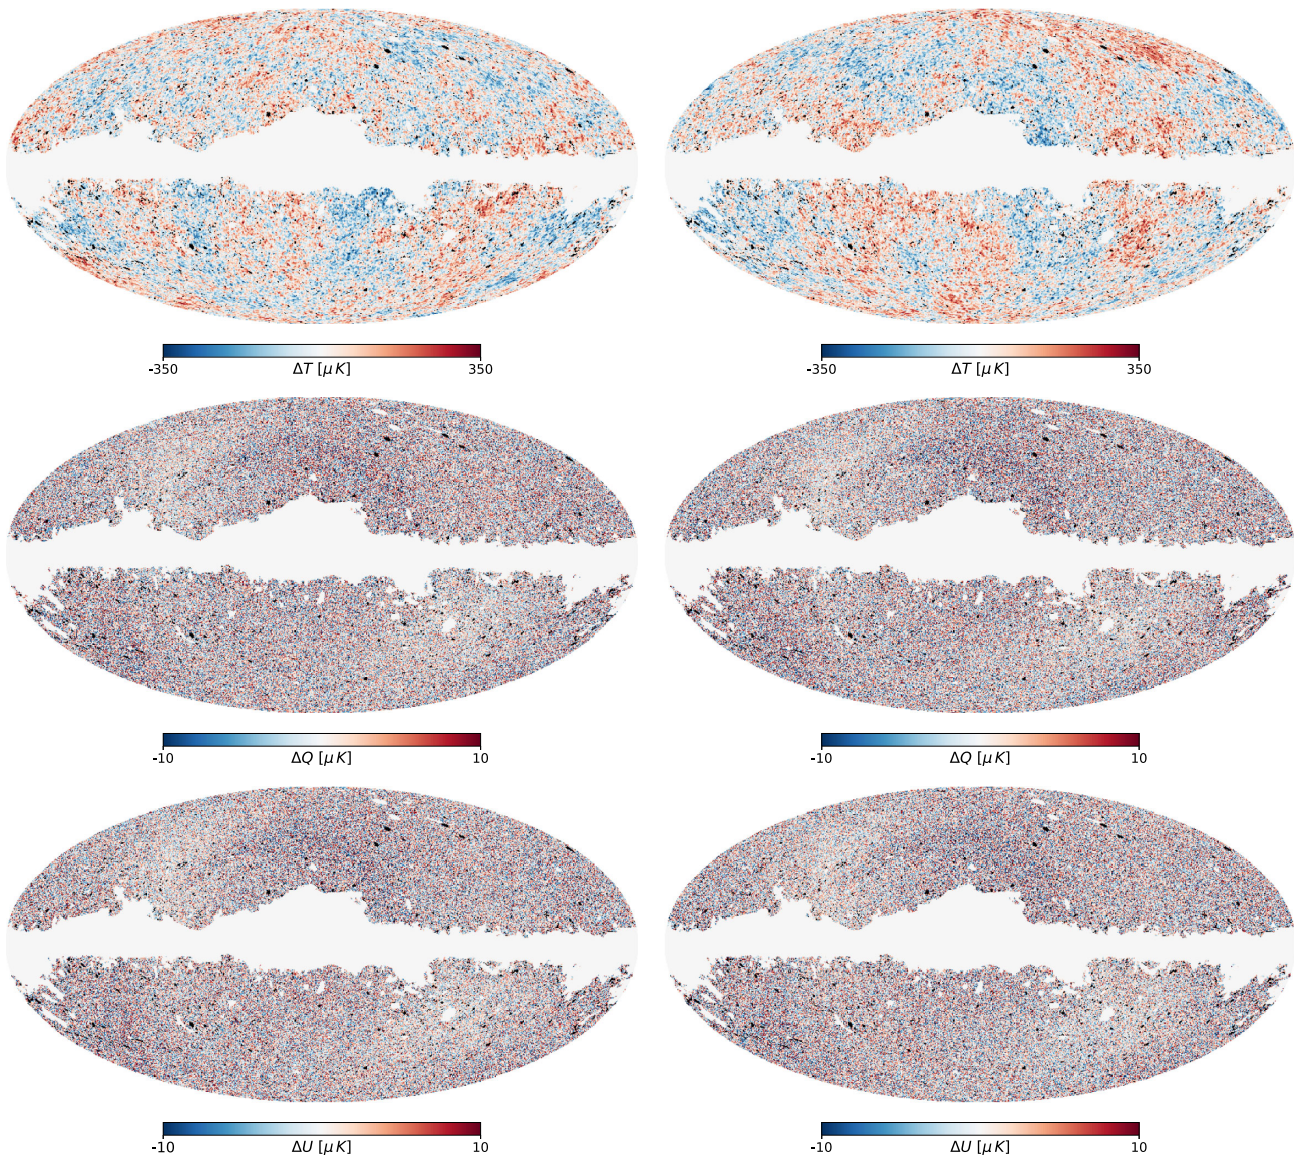


FIG. 1. Comparison between the *Planck* data (left) and a single FFP10 simulation (right). The three plots show T , Q , and U maps (in μK units), all of which are used in the analysis. Black regions show pixels that are inpainted in our analysis, whilst those in white are masked. We note that the FFP10 simulations capture the large-scale (systematic) non-Gaussian features in the polarized maps.

with the addition of an NPIPE polarization transfer function, required to remove spurious effects in the $\ell < 40$ E -mode (and B -mode) power spectrum [38]. The latter is computed from the cross spectra of 600 FFP10 simulations (filtered as above) with the input realizations. Whilst our filtering is not strictly optimal (since we do not fully account for off diagonal harmonic-space correlations induced by the mask, nor translation-invariant noise), such effects are expected to be small except on the largest scales (and do not bias the results) [41]. In Sec. IV, we will additionally make use of Gaussian random field simulations; these are generated as masked Gaussian realizations of the S_{ℓ}^{XY} power spectra.

In Fig. 1, we plot the masked *Planck* temperature and polarization maps, alongside the first FFP10 simulation. The agreement between simulation and data is excellent (though the underlying CMB realization differs); importantly, in both cases, we see non-Gaussian signatures in the polarized maps, exhibiting a quadrupole-like variation in power. This can be attributed to residual foreground effects [38], and could potentially give a spurious signal in the polarized trispectrum. Since this is present in both simulations and data, we can verify our trispectrum pipeline on the FFP10 simulations to check for such effects, and, if necessary (which is not true in this work), subtract their mean value.

III. TRISPECTRUM ESTIMATION

Given the *Planck* and FFP10 simulations, the experimental beam, and the \mathbf{S}^{-1} weighting, we compute the observed trispectra using the POLYBIN code [48], described in [41,42], loosely based on the template approaches of [47,49]. This is a quasioptimal estimator of the full-sky trispectrum, which takes the form

$$t^u(\mathbf{b}, \beta) = \sum_{\mathbf{b}', \beta', u'} \mathcal{F}^{-1, uu'}(\mathbf{b}, \beta; \mathbf{b}', \beta') t^{\text{num}, u'}(\mathbf{b}', \beta'), \quad (3)$$

comprising two contributions: a data-dependent numerator, t^{num} , and a data-independent normalization, \mathcal{F} (both of which can be efficiently computed). Here $t^u(\mathbf{b}, \beta)$ is the trispectrum in ℓ -bins $\mathbf{b} \equiv \{b_1, b_2, b_3, b_4\}$, internal L -bin β and fields $u \equiv \{X_1, X_2, X_3, X_4\}$ (for $X_i \in \{T, E\}$). The numerator is the sum of three contributions:

$$\begin{aligned} t^{\text{num}, u}(\mathbf{b}, \beta) = & \frac{1}{24} \frac{\partial \langle d^i d^j d^k d^l \rangle}{\partial t^u(\mathbf{b}, \beta)} \{ [\mathbf{S}^{-1} d]^i [\mathbf{S}^{-1} d]^j [\mathbf{S}^{-1} d]^k [\mathbf{S}^{-1} d]^l \\ & - 6 \{ [\mathbf{S}^{-1} d]^i [\mathbf{S}^{-1} d]^j \} [\mathbf{S}^{-1} d]^k [\mathbf{S}^{-1} d]^l \\ & + 3 \{ [\mathbf{S}^{-1} d]^i [\mathbf{S}^{-1} d]^j \} \langle [\mathbf{S}^{-1} d]^k [\mathbf{S}^{-1} d]^l \rangle \}, \quad (4) \end{aligned}$$

where i, j, k, l run over all pixels and fields in the dataset, d (assuming Einstein summation). Only the first term is strictly required for the parity-odd estimators (which involves four copies of the data, or $d_i \equiv a_{\ell m}^i$); the other terms remove Gaussian contributions to the parity-even estimator and can reduce variance. The normalization matrix, \mathcal{F} , is set such that the estimator is unbiased for any weighting scheme \mathbf{S}^{-1} , and can ameliorate mask-induced leakage between bins, polarizations, and parities, as well as correlations between T and E , inpainting and beam effects. Explicitly, it is given by

$$\begin{aligned} \mathcal{F}^{uu'}(\mathbf{b}, \beta; \mathbf{b}', \beta') = & \frac{1}{24} \frac{\partial \langle d^i d^j d^k d^l \rangle}{\partial t^u(\mathbf{b}, \beta)} \mathbf{S}_{ii'}^{-1} \mathbf{S}_{jj'}^{-1} \mathbf{S}_{kk'}^{-1} \mathbf{S}_{ll'}^{-1} \\ & \times \frac{\partial \langle d^i d^j d^k d^l \rangle}{\partial t^{u'}(\mathbf{b}', \beta')}. \quad (5) \end{aligned}$$

In the limit of an ideal weighting scheme and a Gaussian dataset (i.e. \mathbf{S}^{-1} equal to the inverse of the pixel covariance, $\langle dd^T \rangle$), the estimators are optimal, such that their covariance is equal to \mathcal{F}^{-1} [41,50]. Although verification of the estimators on parity-violating simulations is beyond the scope of this work, the trispectrum estimators have been extensively tested in [41,42], most notably by assessing their idealized variances and application to Gaussian simulations.

Here, we perform two tests for parity violation. Firstly, we compute the *Planck* trispectra across a broad range of bins and configurations, working at modest ℓ_{max} and $N_{\text{side}} = 256$ (echoing [37]). We compute the trispectra for

each nontrivial combination of T and E -modes ($TTTT$, $TTTE$, $TTEE$, $TETE$, $TEEE$, $EEEE$, with appropriate restrictions on bins to avoid degeneracies), noting that B -modes are not sourced by scalar physics at leading order, and any mask-induced leakage is small (though we consider this explicitly in Appendix B). For this purpose, we use $N_\ell = 7$ linear bins equally spaced in $\ell^{2/5}$ with $\ell \in [2, 518]$ (somewhat coarser than [37] to temper the increase in dimensionality caused by the addition of polarization), such that the signal to noise is roughly constant in each bin (according to the primordial scalings of [51]).⁶ The full trispectrum is specified by four ℓ -bins (for the four sides) and one diagonal L -bin (cf. 1, giving a total of $N_{\text{bin}} = 3130$ bins in the data vector (across all polarizations), or 2386 when the first ℓ -bins are dropped due to reduce leakage from unmeasured bins. By default, we do not include any parity-even modes in the analysis; this assumption is verified in Appendix B. Whilst finer ℓ -bins and larger ℓ_{max} would facilitate (somewhat) higher precision analyses, the corresponding trispectra become significantly more computationally expensive. The above values were chosen to balance these factors.

Secondly, we perform a targeted analysis of a specific gauge-field model (described in [12] and Sec. V), whose signatures primarily lie in collapsed trispectra with $L \lesssim 5$). In this case, we may extend to higher resolution ($N_{\text{side}} = 1024$), since we require far fewer trispectrum bins to characterize the signal (with the number of collapsed bins scaling as ℓ_{max}^2 , rather than ℓ_{max}^5 for general bins). In this case, we use $N_\ell = 7$ linear bins logarithmically spaced in $[2, 2000]$ (chosen to allow easy exploration of the ℓ_{max} -dependence with modest computation resources), but restrict to $L \leq 10$ for the internal leg.⁷ This yields $N_{\text{bin}} = 484$ bins in total, or 342 after subtraction of the smallest ℓ -bin.

In practice, the trispectrum numerator given in (4) is computed from the dataset d via a series of spin-weighted spherical harmonic transforms, as described in [41], with the expectation terms computed via Monte Carlo summation using the FFP10 simulations.⁸ Here, we use 100 simulations for the blind tests, or 10 for the targeted analysis, noting that such terms are sourced only by mask-induced leakage (though can inflate the error bars on large scales). Computation of the disconnected terms is usually rate limiting, and scales as $N_{\text{sim}} N_\ell^2$ for arbitrary trispectra of $N_{\text{sim}} N_\ell$ for collapsed trispectra. Secondly, the normalization matrix, \mathcal{F} , is computed also via Monte Carlo

⁶The bin edges are given by $\{2, 4, 22, 62, 128, 223, 352, 518\}$.

⁷The bin edges are given by $\{2, 5, 10, 30, 100, 300, 1000, 2000\}$.

⁸Note that the error induced by an incorrectly assumed fiducial cosmology starts at second order in $C_\ell^{\text{fid}} - C_\ell^{\text{true}}$ and is heavily suppressed by the restriction to parity-odd modes. The fiducial cosmology also appears in the weighting scheme; this cancels in the narrow-bin limit and cannot induce bias.

methods (in particular, the Girard-Hutchinson stochastic trace estimator [52,53]), making use of $N_{\text{mc}} = 10$ Gaussian random field simulations generated at the fiducial cosmology. This scales linearly with N_{bin} and N_{mc} , though memory is usually the limiting factor. As shown in [41], $N_{\text{mc}} = 10$ realizations facilitates good convergence; for the analysis below, we validate this by computing also trispectra with $N_{\text{mc}} = 5$, which differ by only $\mathcal{O}(0.01\sigma)$ inducing a bias in the parity-violating χ^2 statistic (see Sec. IV) below 2%.

In the above configurations, each blind-test numerator required ≈ 90 CPU hours to compute (involving $\approx 60\,000$ harmonic transforms), whilst the normalization matrix required ≈ 450 CPU hours (and $\approx 80\,000$ harmonic transforms) per realization and 1 TB memory. For the collapsed trispectra, each numerator required ≈ 30 CPU hours (6 000 transforms), whilst each normalization realization required ≈ 2000 CPU hours (15 000 transforms). We note that (a) the numerator run-time could be significantly reduced by using fewer disconnected simulations (whose impact should be small), and (b) the Fisher computation is not optimally parallelized, since memory considerations usually imply exclusive node usage.

IV. RESULTS: BLIND TEST

We begin by performing a ‘‘blind test’’ for parity violation, whereupon we search for a (model-agnostic)

excess in the binned trispectrum relative to the expected noise fluctuations (following the methodology of [37]). To this end, we compute parity-odd trispectra with $\ell \in [2, 518]$ for the *Planck* data, in addition to 600 FFP10 simulations and 1000 Gaussian random field realizations.

A. Simulations and covariances

A necessary ingredient in the estimators is the normalization matrix, \mathcal{F} ; as discussed above, this is equal to the inverse covariance matrix of the trispectrum if the estimator is optimal. In Fig. 2, we test this by comparing the structure of \mathcal{F}^{-1} to the empirical correlation matrix derived from the FFP10 trispectra. Both matrices have a complex structure, which is off diagonal at the $\mathcal{O}(10\%)$ level. This arises due to three main effects: (1) intrinsic correlations between T - and E -modes (which can be large); (2) mask-induced leakage between neighboring bins (which are not necessarily neighboring in our flattened trispectrum representation); (3) intrinsic correlations due to a degeneracy in the trispectrum diagonal definition. Clearly, these effects must be included in any analysis pipeline; here, we find that they are very well described by the inverse normalization matrix, \mathcal{F}^{-1} .

The above discussion motivates us to define the rescaled trispectra:

$$\tau^u(\mathbf{b}, \beta) = [(\mathcal{F}^{1/2})^T t^u(\mathbf{b}, \beta)], \quad (6)$$

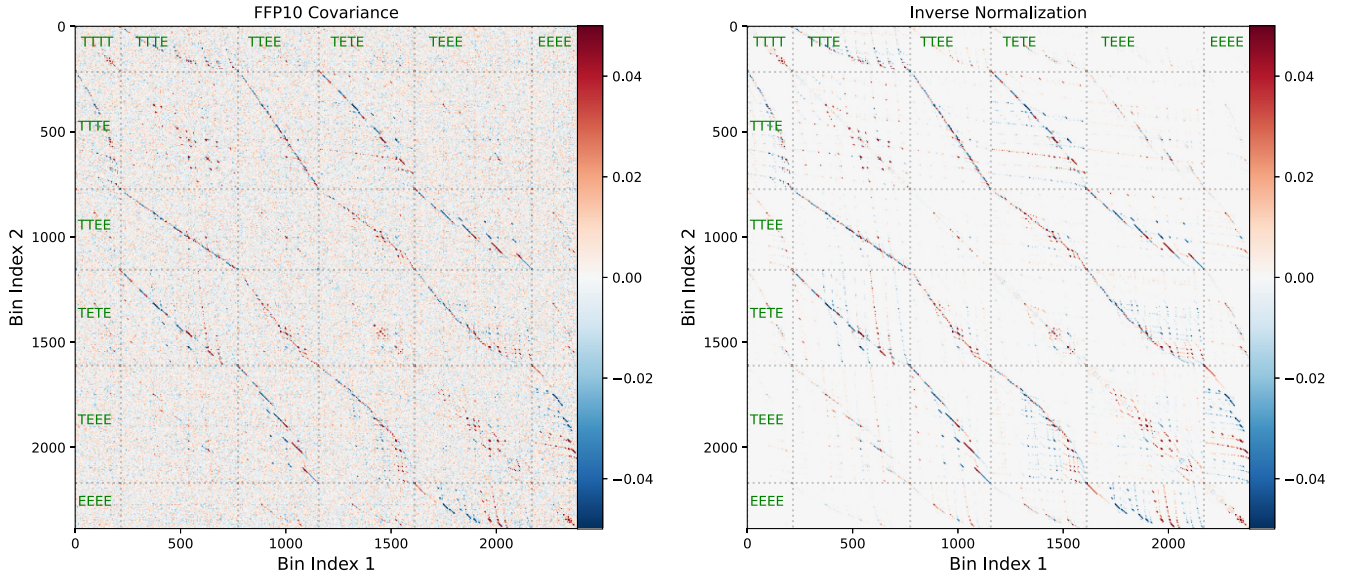


FIG. 2. Comparison of empirical (left) and theoretical (right) correlation matrices for the trispectrum, $t^u(\mathbf{b}, \beta)$. The empirical matrices are computed using 600 FFP10 simulations, whilst the theoretical prediction is given by the inverse trispectrum normalization matrix, itself obtained via Monte Carlo methods, using only the survey mask and weighting scheme. In each panel, we stack all 2386 T - and E -mode trispectrum measurements with $4 \leq \ell < 518$ into one dimension (ordered from low- to high- ℓ), separating out the six nontrivial combinations of fields, which are marked in green. For example, the submatrix in the second column and third row contains the $TTTE \times TTEE$ covariance. Note that different correlators contain different numbers of bins, due to labeling degeneracies (i.e. restricting to $\ell_1 < \ell_2$ for fields $X_1 \neq X_2$). We subtract the (unit) diagonal of the matrix for clarity, and truncate the scale to $\pm 5\%$. The empirical correlation matches the inverse normalization to high precision (up to noise), implying that the latter can be used to decorrelate the measured data, as shown in Fig. 4.




FIG. 3. Rescaled parity-odd trispectrum measured from *Planck* PR4 data, alongside measurements from 600 FFP10 and 1000 GRF simulations for all combinations of T - and E -modes with $4 \leq \ell < 518$. The top panel shows the raw measurements, whilst the bottom shows their simulated variances. The τ -statistic is defined in (6), with weighting chosen to decorrelate the various trispectrum bins. As in Fig. 2, we stack the bins into one dimension, with different polarizations indicated by dotted lines. We find no obvious detection of the trispectrum across the 2386 bins considered herein. The variance of the simulations is larger than unity, implying that the estimators are slightly suboptimal, particularly for polarization. Whilst this does not bias our analysis, it suggests slightly stronger constraints could be wrought with a more advanced weighting scheme.

where $\mathcal{F}^{1/2}$ is the Cholesky factorization of the (binned) normalization matrix \mathcal{F} [defined in (5)].⁹ Noting that the structure of $\mathcal{F}^{-1/2}$ is approximately equal to that of the covariance of t (up to one-dimensional rescalings), this decorrelates the various trispectrum components, which significantly simplifies analysis [37,50]. As a further verification, we plot the correlation matrix of τ in Fig. 4, finding no visible departures from diagonality for both FFP10 and GRF simulations. This is particularly relevant for the FFP10 case, since these simulations contain additional foreground and non-Gaussian contributions.

The final ingredient in our model is the variance of τ . If the trispectrum estimation pipeline is optimal and the statistic is Gaussian distributed, this should be unity (in which case \mathcal{F}^{-1} and the covariance matrix match, instead of just their correlation structures). In the lower panel of Fig. 3 we plot the variance of τ for the two sets of simulations, finding $\tau > 1$ for most bins, with particularly pronounced signatures in bins at low L and for E -modes. This indicates suboptimality (but not bias), and arises from (a) mask- and noise-induced correlations between pixels, (b) violations of the central limit theorem at low- ℓ , and (c) additional variance sourced by the finite number of

⁹Note that the lowest ℓ -bins are removed from the data vector before τ is computed to ensure unbiasedness.

simulations used in the estimator numerator (which are shuffled between measurements). That the signatures are different for the GRF and FFP10 simulations indicates that




FIG. 4. Correlation matrix of the rescaled trispectrum statistic τ , defined by (6). This is obtained using both 600 FFP10 simulations (upper left) and 1000 GRF simulations (bottom right). Unlike the unscaled trispectrum correlations shown in Fig. 2, we find no evidence for off diagonal contamination in the τ covariances, with both matrices consistent with noise. As in Fig. 2, we stack all bins into one dimension and subtract the leading diagonal (which is unity only if the estimator is optimal).

this is impacted also by residual non-Gaussianities induced by lensing, foregrounds, and point sources. We additionally find stronger impacts for polarized fields (with the T -only results matching [37], where relevant), as expected due to the more complex noise and foreground effects therein. We stress that this does not induce bias (provided the FFP10 simulations well represent *Planck*); it simply implies that the eventual sensitivity of our analyses is slightly lower than those that would be obtained if the pixel covariance was perfectly well known.

Putting the above results together, we can define the following compressed statistic used to analyze the data:

$$\chi^2[\hat{\tau}] = \sum_{u, \mathbf{b}, \beta} \frac{(\hat{\tau}^u(\mathbf{b}, \beta) - \tau^{\text{model}, u}(\mathbf{b}, \beta))^2}{\text{var}[\tau^u(\mathbf{b}, \beta)]}, \quad (7)$$

where τ^{model} is some theoretical model (possibly zero), $\hat{\tau}$ is the data, and $\text{var}[\tau]$ is the variance extracted from simulations.¹⁰ If $\hat{\tau}$ is Gaussian distributed, this follows a χ^2 distribution with N_{bins} degrees of freedom; if low- ℓ and low- L modes are important in the analysis, this may be a poor assumption. As such, we will principally use the *empirical* distributions of χ^2 defined from simulations, thus tiptoeing into the realm of frequentist statistics and simulation based inference.

B. Data

In Fig. 3 we plot the rescaled trispectrum measurements from the *Planck* PR4 dataset in addition to those from the simulations. For both sets of simulations, the mean trispectrum is consistent with zero, indicating that our pipeline is not biased. The *Planck* results also show no discernible signal, though the noise makes it difficult to make strong conclusions from visual inspection alone.

To robustly assess the statistical properties of our data, we perform a blind frequentist test for parity violation. Here, we compute the χ^2 statistic of (7) assuming zero theory model, and compare the empirical distributions from simulations to the χ^2 value obtained from *Planck*. This test effectively asks the question: ‘‘Do we find *any* evidence of parity-violating physics in the *Planck* dataset?’’. As such, it is blind to the physical shape of a parity-violating model, and one could obtain stronger constraints on particular models using a dedicated analysis (as in Sec. V).

The results of this test are shown in Fig. 5. We firstly note that the empirical probability density functions (PDFs) are somewhat broader than the expected χ^2 distribution; this suggests that the trispectrum likelihood is non-Gaussian, and implies that a classical χ^2 -analysis could yield biased results, whence noise fluctuations in the tails are interpreted




FIG. 5. Probability of parity violation in the large-scale ($4 \leq \ell < 518$) *Planck* PR4 dataset. We compare the χ^2 statistic of (7) measured in *Planck* to empirical distributions from FFP10 and GRF simulations (red and green respectively). We find excellent agreement between *Planck* results and the simulation histograms, with frequentist ranks and χ^2 probabilities to exceed shown in the title. Both empirical distributions exhibit slightly heavier tails than the theoretical distribution, implying that classical χ^2 -analyses are insufficient for modeling the trispectrum noise properties. The frequentist detection probability for *Planck* corresponds to -0.4σ .

as high-significance detections (which may partially explain the results of [33]). The FFP10 and GRF distributions are generally consistent; this implies that foreground and lensing effects (present only in FFP10) are not leading to additional non-Gaussianity (though we note that the two sets of simulations have different variances, and modeling the FFP10 trispectrum with the GRF variance would lead to biased results). Turning to the *Planck* data, we find a χ^2 value highly consistent with both empirical distributions, with a rank of 201/600 for FFP10 (equivalent to -0.4σ) or 340/1000 for the GRFs.¹¹ Given that both sets of simulations are parity conserving, this test leads to the following conclusion: *we find no evidence of scalar parity violation in the large-scale ($4 \leq \ell < 518$) *Planck* polarized dataset.*

Although the above analysis has returned a null result, it is important to assess the various assumptions in the measurement. To this end, we consider the following variations of the analysis in Appendix B:

- (1) *Inclusion of B-modes*: This accounts for mask-induced leakage between E - and B -modes, by additionally estimating trispectra containing a single B -mode (e.g., $TTTB$, $TETB$), then discarding them after application of the inverse normalization matrix.

¹⁰We use FFP10 simulations for this in general, though we analyze the GRF distribution with the GRF variance, since these clearly differ for E -modes (Fig. 3).

¹¹If using only the GRF simulations [for the empirical distribution and the τ variance entering (6)], the *Planck* data has a rank of 906/1000, equivalent to 1.3σ .

- (2) *Inclusion of parity-even trispectra*: This removes any leakage between parity-odd and parity-even modes induced by the mask, analogous to the B -mode case.
- (3) *Removal of four-field terms*: This assesses the impact of the disconnected trispectrum terms by computing only the full trispectrum, $\propto \langle \prod_{i=1}^4 a_{\ell_i m_i}^{X_i} \rangle$, which does not require any Monte Carlo estimation for the trispectrum numerator. In the ideal limit, such terms do not contribute to parity-odd spectra (though they may enhance variances).
- (4) *Removal of polarization*: This assesses the impact of the E -mode spectra, by restricting to $TTTT$ -trispectra only. This additionally validates the results of [37] with updated *Planck* PR4 data.

The first and second procedures are somewhat more expensive than the fiducial approach; thus we restrict to 100 simulations in both analyses. In all cases, we find no evidence for parity violation.

V. RESULTS: GAUGE MODEL TEST

In [37], a variety of models were compared to the (broadly binned) *Planck* parity-odd temperature trispectra to yield constraints on their amplitudes. Here, we take the most promising model (“detected” at $\approx 2\sigma$ previously) and perform a dedicated polarized trispectrum analysis to yield tight constraints on its amplitude. We utilize the model described in [12], which invokes the coupling between a pseudoscalar field ϕ and a gauge field A_μ to yield parity-odd trispectra via a tree-level interaction sourced by a Chern-Simons term in the inflationary Lagrangian, involving the vacuum expectation value of the gauge field. In terms of the curvature perturbation ζ , the trispectrum is a special case of the general collapsed form introduced in [12]

$$\begin{aligned}
 T_\zeta(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3, \mathbf{k}_4) = & -i \sum_{n \geq 0} d_n^{\text{odd}} [L_n(\hat{\mathbf{k}}_1 \cdot \hat{\mathbf{k}}_3) \\
 & + (-1)^n L_n(\hat{\mathbf{k}}_1 \cdot \hat{\mathbf{s}}) + L_n(\hat{\mathbf{k}}_3 \cdot \hat{\mathbf{s}})] \\
 & \times (\hat{\mathbf{k}}_1 \times \hat{\mathbf{k}}_3 \cdot \hat{\mathbf{s}}) P_\zeta(k_1) P_\zeta(k_3) P_\zeta(s) \\
 & + 23 \text{ perms,} \quad (8)
 \end{aligned}$$

where $\mathbf{s} \equiv \mathbf{k}_1 + \mathbf{k}_2$, P_ζ is the primordial power spectrum and L_n is a Legendre polynomial of integer order. In the gauge field scenario, the amplitude of the model is set by $A_{\text{gauge}} = -d_0^{\text{odd}} = d_1^{\text{odd}}/3$ (equal to A_{CS} of [37]), which is proportional to the fractional energy density in the gauge field’s vacuum expectation value. For generality, we will additionally consider the joint constraints on d_0^{odd} and d_1^{odd} below, though we note that the $n = 1$ term is strongly dominant in the above expansion [cf. [12]].

Further discussion of the model can be found in [12], as well as its application to large-scale structure and CMB

in [21,32,37] respectively. Here, we extend the model to the polarized sector using a sum-separable calculation described in Appendix A. Whilst the mathematics is a little involved, the end product is simple: a set of binned trispectrum amplitudes $t^{\text{th,u}}(\mathbf{b}, \beta)$ that can be rescaled to τ measurements via (6). Since the model is dominated by configurations with low- L (analogous to the τ_{NL} parity-even template), we restrict our analysis to collapsed tetrahedra, which allows for an extension to higher ℓ_{max} in reasonable computation time, as discussed in Sec. III.

In Fig. 6, we plot the parity-odd trispectrum predicted by the first two terms in the expansion of (8), calculated as described in Appendix A, alongside the empirical trispectrum variance from the FFP10 simulations. Being a projection of high-dimensional data, the figure is difficult to interpret: however, it is clear that the trispectrum amplitudes are largest at low- ℓ (particularly at low L), and generally suppressed in E -modes compared to T -modes. The latter occurs due to the smaller transfer function, and echoes the lower signal to noise found in the *Planck* EE power spectrum compared to TT . In comparison to the variance, it is clear that the strongest signatures will be found in $TTTT$ and $TTTE$ trispectra, particularly at high- ℓ (where the variance is suppressed). Finally, the $n = 0$ and $n = 1$ shapes (which differ in their angular complexity), have markedly different signatures, with the $n = 1$ case (which dominates the gauge field model) having significantly larger amplitude, since it is enhanced in the collapsed limit [cf. [12]].

Using (7), we can compare the observed and theoretical trispectra, and thus place constraints on the gauge field model amplitude, A_{gauge} (or the model-agnostic amplitude d_n^{odd}). For this purpose, we assume a Gaussian posterior for the amplitude parameters with a flat (pseudo)prior of infinite extent; this is appropriate via the central limit theorem. The results of this analysis are shown in Fig. 7. From the FFP10 simulation results, we verify that our pipeline is unbiased in all cases, with all the inferred value of A_{gauge} consistent with zero. From the *Planck* data, we also find no detection of A_{gauge} , with a 1σ contour $A_{\text{gauge}} = (-0.9 \pm 3.3) \times 10^3$ for the fiducial analysis. This can be translated into a constraint on the energy density of the gauge field relative to the inflaton (assuming fiducial parameters with coupling $\gamma = 1$): $\rho_{\text{gauge}}/\rho_\phi = (-0.7 \pm 2.6) \times 10^{-19}$. As shown in Fig. 8, we may similarly constrain the d_n^{odd} parameters; the *Planck* data yield $d_0^{\text{odd}} = (-0.9 \pm 1.1) \times 10^9$, $d_1^{\text{odd}} = (-3.1 \pm 9.8) \times 10^3$, with a much stronger bound for the $n = 1$ mode, as expected. The conclusion is clear; *we report no evidence for the parity-odd trispectrum models*.

In Figs. 7–9, we additionally assess the dependence of our constraints on a number of modeling choices. Firstly, we consider removing the E -mode data. This inflates constraints by a factor $\approx 3\times$, indicating that large-scale




FIG. 6. Theoretical models for the CMB T - and E -mode trispectra, alongside the empirical trispectrum error extracted from 100 FFP10 simulations. We show the $n = 0$ and $n = 1$ templates given in (8), the second of which well approximates the gauge field model. In contrast to previous figures, results are shown for $5 \leq \ell < 2000$, with internal $2 \leq L < 10$, where the model is large. The fiducial amplitudes of the model are set to $d_n^{\text{odd}} = 10^4$ for visualization (corresponding to $A_{\text{gauge}} = 10^4/3$), and we take the absolute value of the model correlators. Comparing theory and error, we expect that the constraints will be dominated by $TTTT$ and $TTTE$ trispectra on small scales.




FIG. 7. Constraints on the amplitude of the gauge-field model [12] for the parity-violating trispectrum. We show the posterior for both the *Planck* PR4 data (up to $\ell = 2000$) and the mean of 100 FFP10 simulations. The left panel shows the fiducial analysis with internal $L \geq 2$, whilst the central and right panels respectively show the results of removing E -modes and restricting to $L \geq 5$. The fiducial results find $A_{\text{gauge}} = (-0.9 \pm 3.3) \times 10^3$ (*Planck*) and $A_{\text{gauge}} = (-0.5 \pm 3.3) \times 10^3$ (FFP10), with no detection of parity violation. Constraints degrade by a factor ≈ 3 with the excision of E modes, or by a factor ≈ 35 with the removal of low- L configurations.

polarization data can be of great use when constraining primordial models. For the *Planck* data, the T -only results show a slight ($\approx 2\sigma$) preference for $A_{\text{gauge}} > 0$ and $d_1^{\text{odd}} > 0$ (Figs. 7 and 8); this is here identified as a statistical fluctuation, given that (a) the FFP10 results are unbiased and (b) the deviation vanishes when E -modes are included. Secondly, we consider the removal of the first L -bin, giving

$L \geq 5$. This degrades the constraints by a factor of ≈ 35 , implying that the model is dominated by the lowest L -modes, as discussed in [12]. This implies that the measurement would be highly sensitive to large-scale systematics, such as foreground-induced quadrupolar anisotropies in the power spectrum, analogously to τ_{NL} constraints. Thirdly, we repeat the analysis without removing the




FIG. 8. Constraints on the general parity-odd trispectrum given in (8) from the mean of 100 FFP10 simulations (blue) and *Planck* data (red) up to $\ell_{\max} = 2000$, as in Fig. 7. We show two sets of contours: T -only (the wider and lighter contours), and $T + E$ (the darker and narrower contours). The best-fit *Planck* constraints are $d_0^{\text{odd}} = (-0.9 \pm 1.1) \times 10^9$, $d_1^{\text{odd}} = (-3.1 \pm 9.8) \times 10^3$ from temperature and polarization, or $d_0^{\text{odd}} = (-7.9 \pm 3.1) \times 10^9$, $d_1^{\text{odd}} = (4.5 \pm 3.3) \times 10^4$ from temperature alone.

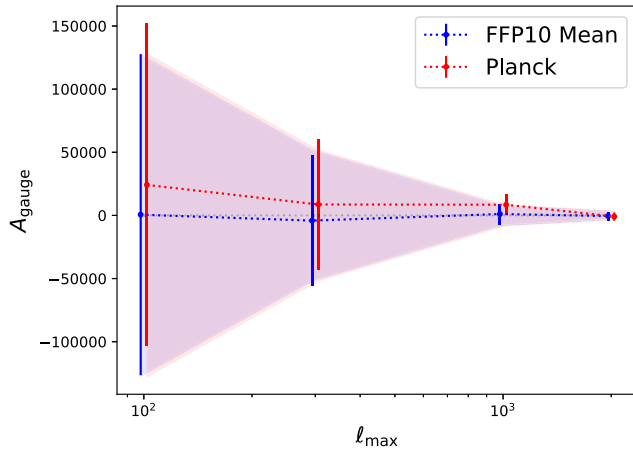


FIG. 9. Dependence of the gauge-model constraints shown in Fig. 7 on scale cuts, via the maximum multipole ℓ_{\max} . We show the 1σ error bar both for the mean of 100 FFP10 simulations (blue) and *Planck* PR4 data (red). We include both T - and E -modes in all analyses and use internal $2 \leq L < 10$. We find strong dependence on ℓ_{\max} , though the scaling tempers slightly beyond $\ell_{\max} \gtrsim 1000$, likely due to projection effects.

disconnected trispectrum contributions (as discussed in Sec. IV); this is found to bias the mean-of-FFP10 results at 1σ , thus we conclude that such terms are important (in the presence of inhomogeneous noise and masking). Finally, we consider the scale dependence for the external trispectrum legs in Fig. 9, whence we constrain A_{gauge} using subsets of

the full data vector. Whilst this approach does not quite emulate a complete analysis at lower ℓ_{\max} (which would feature a finer binning), it nevertheless indicates a strong dependence of $\sigma(A_{\text{gauge}})$ on ℓ_{\max} . Here, we find this to be somewhat slower than quadratic, despite the predictions of [12,51] (for a τ_{NL} -like diagram). This is likely due to a collection of effects: (a) nonlinearity in the polarization covariance on small scales (captured in our simulated-based variances), (b) damping and projection effects at $\ell \gtrsim 1000$, and (c) suboptimal binning strategies (chosen based on computational considerations). Using finer ℓ -bins may lead to slightly tighter bounds, though, in the absence of a definitive target for A_{gauge} or d_n^{odd} , it is unclear to what extent this is useful.

VI. DISCUSSION

In the above sections, we have presented novel constraints on the parity-odd sector of the Universe using *Planck* temperature and polarization data, facilitated by efficient new trispectrum estimators [41,42]. It is now our duty to place these in context: what does this imply for observations and models? How can the analysis be extended? What is the outlook for the future?.

The blind analysis performed in Sec. IV places competitive constraints on scalar parity-violating phenomena, both in the primordial Universe, and at late times, through CMB secondary effects. Importantly, these results do not depend on a physical model: any process with sufficiently large amplitude should show up as an excess signal in the simulation-based χ^2 -analysis. That said, our analysis is limited to comparatively large scales ($\ell < 518$), due to the computational burden, and we cannot rule out the possibility of signals lurking at higher- ℓ (for example in CMB secondaries, such as lensing).

It is interesting to consider the strength of such constraints in relation to previous CMB temperature and galaxy analyses [32,33,37]. As shown in [37], the *Planck* dataset provides far stronger constraints on scale-invariant primordial models than the BOSS galaxy survey, since it contains $\approx 250\times$ more primordial modes. As noted in [37], this strongly implies that the signatures of parity violation reported in [32,33] are not primordial in origin. The addition of E -modes can only boost the signal to noise: the precise gain, however, will depend strongly on the types of primordial model considered. At low $\ell \lesssim 5$, the number of modes in $T + E$ analyses is $\approx 80\%$ larger than those in T -alone (quantified as in [37], cf. [54] for galaxy surveys), though the ratio drops significantly at high- ℓ , whence polarization noise dominates. As such, models with important contributions at low- ℓ can be expected to benefit significantly from the addition of E -modes. The situation, in practice, is a little more complex, since primordial physics affects T - and E -modes in different ways, with different sensitivity to noise and projection effects. As such,

the true gain can be obtained only by testing individual models (or through Fisher forecasts).

In Sec. V, we considered one such model: the $U(1)$ gauge-field scenario of [37]. In this case, inclusion of E -mode data significantly bolstered the constraints (by a factor of ≈ 3), with a 1σ bound of $A_{\text{gauge}} = (-0.9 \pm 3.3) \times 10^3$, or $\rho_{\text{gauge}}/\rho_\phi = (-0.7 \pm 2.6) \times 10^{-19}$ for the fractional gauge field energy density. This is in accordance with the Fisher forecasts of [12, Fig. 4], but several orders of magnitude stronger than the previously reported bounds of $A_{\text{gauge}} = (-2.8 \pm 1.3) \times 10^7$, from *Planck* temperature [37]. The huge improvement over the former work is due to a number of factors: (a) extension to much smaller scales (from $\ell_{\text{max}} \approx 500$ to 2000), (b) inclusion of the lowest L -modes (down to $L = 2$), (c) addition of E -mode data, and (d) improved calculation of the theoretical templates, with less noise bias. The large-scale structure results from [32] appear slightly more optimistic than those found herein; as noted in [37], this is erroneous, since the former works did not account for mode decorrelation from nonlinear effects, thus heavily overestimated the primordial physics content of the galaxy statistics (by around 2 orders of magnitude, based on mode counting). Finally, we note that, whilst we have only focused on one model in this study, the conclusions are more general: any parity-violating model of interest can now be constrained using the *Planck* T - and E -mode data. Though these are already indirectly bound from the blind analyses, targeted analyses will always yield stronger constraints due to optimal weighting of the signal.

The next decade will see considerable improvements in the volume and quality of CMB data; thus there is significant hope for obtaining stronger bounds on the above models and other types of new physics. As discussed in Sec. II, gains on primordial physics will be driven by polarization rather than temperature, due to the cosmic

variance limit, and proliferation of secondary effects in temperature such as the kinetic Sunyaev-Zel'dovich signal. As demonstrated in [51], there is much to gain from larger ℓ_{max} (which is currently ≈ 600 in polarization, due to noise limitations), though pushing beyond $\ell_{\text{max}} \approx 2000$ will be hampered by projection effects. Collapsed and squeezed models show particularly significant dependence on ℓ_{max} , with predicted scalings up to $(S/N)^2 \sim \ell_{\text{max}}^4$ depending on the precise inflationary model considered. Even though we have detected nothing in this work, there is strong hope for the future!

The supporting data for this Letter are openly available from [55].

ACKNOWLEDGMENTS

We thank Will Coulton, Adri Duivenvoorden, and Masahiro Takada for insightful discussions. We are additionally grateful to the anonymous referee for insightful comments. O. H. E. P. is a Junior Fellow supported by the Simons Society of Fellows. M. S. is supported by JSPS KAKENHI Grants No. JP20H05859 and No. JP23K03390. O. H. E. P. thanks IPMU for hosting a visit during which this work was started, as well as Llawrence the Llama for spiritual guidance.

APPENDIX A: GAUGE-FIELD MODEL

Here, we outline the derivation and computation of the gauge-field trispectrum model described in [12]. This follows an analogous derivation in [41], but is extended to include E -mode polarization, and a more efficient factorization scheme. We begin with the relation of the CMB harmonic coefficients and the primordial curvature trispectrum, T_ζ ,

$$\left\langle \prod_{i=1}^4 a_{\ell_i m_i}^{X_i} \right\rangle_c = (4\pi)^4 \left[\prod_{i=1}^4 i^{\ell_i} \int_{\mathbf{k}_i} \mathcal{T}_{\ell_i}^{X_i}(k_i) Y_{\ell_i m_i}^*(\hat{\mathbf{k}}_i) \right] T_\zeta(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3, \mathbf{k}_4) (2\pi)^3 \delta_{\text{D}}(\mathbf{k}_{1234}), \quad (\text{A1})$$

for curvature transfer function \mathcal{T}_ℓ^X , with only $\mathcal{T}_\ell^B = 0$ at leading order. Following [41], we can write the unbinned trispectrum [the rhs of (A1)] as

$$\begin{aligned} T_{m_1 m_2 m_3 m_4}^{\ell_1 \ell_2 \ell_3 \ell_4, u} &= -\sqrt{2} i^{\ell_{1234}} A_{\text{gauge}} (\Delta_\zeta^2)^3 (4\pi)^{11/2} \sum_{L_1 L_2 L_3 L_4 L' L''} i^{L_{1234} - L' + L''} C_{L_1 L_2 L_3 L_4 L' L''} \begin{pmatrix} L_1 & L_2 & L' \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} L_3 & L_4 & L'' \\ 0 & 0 & 0 \end{pmatrix} \\ &\times \int x^2 dx \int y^2 dy h_{-3}^{\ell_1 L_1, u_1}(x) h_0^{\ell_2 L_2, u_2}(x) h_{-3}^{\ell_3 L_3, u_3}(y) h_0^{\ell_4 L_4, u_4}(y) f_{L' L''}(x, y) \\ &\times \int d\hat{\mathbf{s}} \prod_{i=1}^4 \left[\int d\hat{\mathbf{k}}_i Y_{\ell_i m_i}^*(\hat{\mathbf{k}}_i) \right] \mathcal{P}_{L_1 L_2 L'}(\hat{\mathbf{k}}_1, \hat{\mathbf{k}}_2, \hat{\mathbf{s}}) \mathcal{P}_{L_3 L_4 L''}(\hat{\mathbf{k}}_3, \hat{\mathbf{k}}_4, \hat{\mathbf{s}}) \\ &\times \sum_{\lambda_1 \lambda_3 \lambda} A_{\lambda_1 \lambda_3 \lambda} \mathcal{P}_{\lambda_1 \lambda_3 \lambda}(\hat{\mathbf{k}}_1, \hat{\mathbf{k}}_3, \hat{\mathbf{s}}) + 23 \text{ perms}, \end{aligned} \quad (\text{A2})$$

where u is a quadruplet of fields. This uses the definitions

$$h_{\alpha}^{\ell L, X}(x) = \int \frac{k^2 dk}{2\pi^2} k^{\alpha} T_{\ell}^X(k) j_L(kx), \quad f_{LL'}(x, y) = \int \frac{s^2 ds}{2\pi^2} \frac{j_L(sx) j_{L'}(sy)}{s^3} \quad (\text{A3})$$

with $A_{111} = 1, A_{221} = -A_{212} = A_{122} = 1/\sqrt{5}$ and zero else.¹² Furthermore, $\mathcal{C}_{\ell_1 \dots \ell_n} = \sqrt{(2\ell_1 + 1) \dots (2\ell_n + 1)}$ and \mathcal{P} are the isotropic basis functions described in [56], trivially related to the tripolar spherical harmonics [57]. Performing the angular integral as for the scalar case, we find

$$\begin{aligned} i_{\ell_1 \ell_2 \ell_3 \ell_4}^u(L) \begin{pmatrix} \ell_1 & \ell_2 & L \\ -1 & -1 & 2 \end{pmatrix} \begin{pmatrix} \ell_3 & \ell_4 & L \\ -1 & -1 & 2 \end{pmatrix} &= -\sqrt{2} A_{\text{gauge}} (\Delta_{\xi}^2)^3 (4\pi)^5 \sum_{L_1 L_3 L' L''} i^{-\ell_1 + \ell_3 - L_1 + L_3 + L' + L''} \mathcal{C}_{L_1 L_3 L' L''}^2 \\ &\times \sum_{\lambda_1 \lambda_3 \lambda} A_{\lambda_1 \lambda_3 \lambda} \mathcal{C}_{\lambda_1 \lambda_3 \lambda} \left\{ \begin{matrix} L & \ell_1 & \ell_2 \\ L_1 & L' & \lambda_1 \end{matrix} \right\} \left\{ \begin{matrix} \lambda & \lambda_1 & \lambda_3 \\ L & L'' & L' \end{matrix} \right\} \left\{ \begin{matrix} \ell_3 & \ell_4 & L \\ L'' & \lambda_3 & L_3 \end{matrix} \right\} \\ &\times \int x^2 dx \int y^2 dy h_{-3}^{\ell_1 L_1, u_1}(x) h_0^{\ell_2 \ell_2, u_2}(x) h_{-3}^{\ell_3 L_3, u_3}(y) h_0^{\ell_4 \ell_4, u_4}(y) f_{L' L''}(x, y) \\ &\times \begin{pmatrix} L_1 & \ell_2 & L' \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} L_3 & \ell_4 & L'' \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} L_1 & \lambda_1 & \ell_1 \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} L' & \lambda & L'' \\ 0 & 0 & 0 \end{pmatrix} \\ &\times \begin{pmatrix} \lambda_3 & L_3 & \ell_3 \\ 0 & 0 & 0 \end{pmatrix}, \end{aligned} \quad (\text{A4})$$

which can be partially factorized into $(\ell_1, \ell_2, u_1, u_2)$ and $(\ell_3, \ell_4, u_3, u_4)$ contributions, linked by integrals and $6j$ symbols.

To compare theory and data, we must appropriately bin the theory model. Following [41], the binned trispectra are defined as

$$\begin{aligned} \hat{i}^u(\mathbf{b}, \beta) &\propto \frac{8}{\Delta_4^u(\mathbf{b})} \sum_{\ell_i L} \Theta_{\ell_1}(b_1) \dots \Theta_{\ell_4}(b_4) \Theta_L(\beta) (2\ell_1 + 1) \dots (2\ell_4 + 1) (2L + 1) \\ &\times \begin{pmatrix} \ell_1 & \ell_2 & L \\ -1 & -1 & 2 \end{pmatrix}^2 \begin{pmatrix} \ell_3 & \ell_4 & L \\ -1 & -1 & 2 \end{pmatrix}^2 \sum_{u'} S_{\ell_1}^{-1, u_1 u'_1} \dots S_{\ell_4}^{-1, u_4 u'_4} i_{\ell_1 \ell_2 \ell_3 \ell_4}^{u'}(L), \end{aligned} \quad (\text{A5})$$

dropping constant factors, and asserting that the spectra is parity odd. Here, the u' sum is over all 2^4 quadruplets of fields (not just the ordered pairs appearing in the trispectra of interest), and the forms are normalized by the diagonal-in- β matrix

$$\begin{aligned} \mathcal{N}^{uu'}(\mathbf{b}, \mathbf{b}', \beta) &\equiv \frac{1}{\Delta_4^u(\mathbf{b}) \Delta_4^{u'}(\mathbf{b}')} \sum_{\ell_i L} \Theta_{\ell_1}(b_1) \dots \Theta_{\ell_4}(b_4) \Theta_L(\beta) (2\ell_1 + 1) \dots (2\ell_4 + 1) (2L + 1) \\ &\times \begin{pmatrix} \ell_1 & \ell_2 & L \\ -1 & -1 & 2 \end{pmatrix}^2 \begin{pmatrix} \ell_3 & \ell_4 & L \\ -1 & -1 & 2 \end{pmatrix}^2 \left[S_{\ell_1}^{-1, u_1 u'_1} \dots S_{\ell_4}^{-1, u_4 u'_4} \delta_{b_1 b'_1}^K \dots \delta_{b_4 b'_4}^K + 7 \text{ perms} \right]. \end{aligned} \quad (\text{A6})$$

This correlates only pairs of bins where \mathbf{b}' is a permutation of \mathbf{b} . The normalization is high dimensional but can be efficiently computed by factorizing:

$$\begin{aligned} \mathcal{N}^{uu'}(\mathbf{b}, \mathbf{b}', \beta) &\equiv \frac{1}{\Delta_4^u(\mathbf{b}) \Delta_4^{u'}(\mathbf{b}')} \sum_p \sum_L \Theta_L(\beta) (2L + 1) \\ &\times \left(\sum_{\ell_1 \ell_2} \Theta_{\ell_1}(b_1) \Theta_{\ell_2}(b_2) (2\ell_1 + 1) (2\ell_2 + 1) \begin{pmatrix} \ell_1 & \ell_2 & L \\ -1 & -1 & 2 \end{pmatrix}^2 S_{\ell_1}^{-1, u_1 u'_{p_1}} S_{\ell_2}^{-1, u_2 u'_{p_2}} \delta_{b_1 b'_{p_1}}^K \delta_{b_2 b'_{p_2}}^K \right) \\ &\times \left(\sum_{\ell_3 \ell_4} \Theta_{\ell_3}(b_3) \Theta_{\ell_4}(b_4) (2\ell_3 + 1) (2\ell_4 + 1) \begin{pmatrix} \ell_3 & \ell_4 & L \\ -1 & -1 & 2 \end{pmatrix}^2 S_{\ell_3}^{-1, u_3 u'_{p_3}} S_{\ell_4}^{-1, u_4 u'_{p_4}} \delta_{b_3 b'_{p_3}}^K \delta_{b_4 b'_{p_4}}^K \right), \end{aligned} \quad (\text{A7})$$

¹²The general trispectrum model of (8) can be computed via the replacements $A_{\text{gauge}} A_{111} \rightarrow -d_0^{\text{odd}}$, and $A_{\text{gauge}} A_{221} = -A_{\text{gauge}} A_{121} = A_{\text{gauge}} A_{122} = d_1^{\text{odd}} / (3\sqrt{5})$.

where p are the eight permutations above of $\{1, 2, 3, 4\}$.¹³ Inserting (A4), the binned trispectrum numerator can be written in factorized form:

$$\begin{aligned}
t^u(\mathbf{b}, \beta) &\propto -\sqrt{2}A_{\text{gauge}}(\Delta_\zeta^2)^3(4\pi)^5 \frac{8}{\Delta_4^u(\mathbf{b})} \sum_{\lambda_1\lambda_3\lambda} (-1)^{\lambda_3} A_{\lambda_1\lambda_3\lambda} C_{\lambda_1\lambda_3\lambda} \\
&\times \sum_{L'L''} \Theta_L(\beta)(2L+1)(2L'+1)(2L''+1) \begin{pmatrix} L' & \lambda & L'' \\ 0 & 0 & 0 \end{pmatrix} \begin{Bmatrix} \lambda & \lambda_1 & \lambda_3 \\ L & L'' & L' \end{Bmatrix} \\
&\times \frac{1}{2} \int \frac{s^2 ds}{2\pi^2 s^3} \left[f_{LL'\lambda_1}^{u_1 u_2}(b_1, b_2; s) f_{LL''\lambda_3}^{u_3 u_4}(b_3, b_4; s) - \bar{f}_{LL'\lambda_1}^{u_1 u_2}(b_1, b_2; s) \bar{f}_{LL''\lambda_3}^{u_3 u_4}(b_3, b_4; s) \right], \quad (\text{A8})
\end{aligned}$$

defining

$$\begin{aligned}
f_{LL'\lambda_1}^{u_1 u_2}(b_1, b_2; s) &= \sum_{\ell_1 \ell_2 L_1} \Theta_{\ell_1}(b_1) \Theta_{\ell_2}(b_2) (2\ell_1+1)(2\ell_2+1) i^{-\ell_1-L_1+L'} (2L_1+1) \\
&\times \begin{pmatrix} L_1 & \ell_2 & L' \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} L_1 & \lambda_1 & \ell_1 \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} \ell_1 & \ell_2 & L \\ -1 & -1 & 2 \end{pmatrix} \begin{Bmatrix} L & \ell_1 & \ell_2 \\ L_1 & L' & \lambda_1 \end{Bmatrix} \\
&\times \int x^2 dx \tilde{h}_{-3}^{\ell_1 L_1, u_1}(x) \tilde{h}_0^{\ell_2 \ell_2, u_2}(x) j_{L'}(sx), \quad (\text{A9})
\end{aligned}$$

where $\tilde{h}_\alpha^{\ell L, u}(y) \equiv \sum_{u'} S_{uu'}^{-1} h_\alpha^{\ell L, u'}(y)$. \bar{f} is analogous but with an extra factor of $(-1)^{\ell_1+\ell_2}$ or $(-1)^{\ell_3+\ell_4}$ (to impose an odd ℓ_{1234} sum). In practice, the various integrals (appearing in h , f and t) are evaluated using numerical quadrature, with linearly spaced arrays in $\{k_i/\ell_i, x, s\}$, with a total of $\{5000, 2000, 1000\}$ points in each (noting that the h integrals can all be precomputed; thus a dense k -grid is possible). Various modifications to the finite integration grids were considered to ensure that any systematic errors induced are small.

APPENDIX B: CONSISTENCY CHECKS

In this Appendix, we present various consistency checks of the blind analysis. In each case, we will reproduce Fig. 5 with the modified assumptions as described in the last paragraph of Sec. IV. For consistency, we compare the B -mode and even-parity runs against fiducial analyses using 100 GRF and FFP10 simulations only. We note that the associated increased noise in the variance of τ leads to a slight upward shift in the empirical χ^2 posteriors in all cases, with a fiducial *Planck* parity-violation probability of 0.06σ .

1. Adding B -modes

As described in Sec. IV, we test for leakage between E - and B -modes (and lensing effects) by additionally including *TTTB*, *TTEB*, *TETB*, *TEEB*, and *EEEE* trispectra in our measurement pipeline (leading to a total of 8360 bins),

¹³This approach was not considered in [41], leading to the prohibitive $\mathcal{O}(\ell_{\text{max}}^5)$ scaling claimed previously.

then discarding such configurations after normalization by \mathcal{F}^{-1} . Assuming the leakage to be small, such spectra (with one $E \rightarrow B$ replacement) dominate the signal. This roughly doubles the computation cost of trispectra, thus is avoided in the main analysis. From Fig. 10(a), we find no significant changes to our main results, except for slight fluctuations in the χ^2 of any given simulation, due to the associated modifications to the trispectrum noise. We find no evidence for parity violation, with the *Planck* data consistent with FFP10 at -0.01σ (from the rank tests).

2. Adding parity-even modes

Leakage between parity-odd and parity-even modes (odd and even $\ell_1 + \ell_2 + \ell_3 + \ell_4$) is assessed in a similar manner by jointly estimating all trispectrum configurations (a total of 6260), then removing the latter after the (decorrelating) normalization matrix is applied (which is block diagonal in the ideal case). From Fig. 10(b), we again find no significant change to the posterior, with smaller noise fluctuations in this case. The *Planck* data have an overall probability of -0.03σ , and we find no evidence for parity violation.

3. Removing disconnected corrections

In Fig. 10(c) we assess the impact of the disconnected terms in the trispectrum numerator. In the ideal case, these will not modify the signal, but may change the covariance of certain bins (specifically, those with two or four equal ℓ -bins). Since this study does not require new analysis, we can use all the previously computed fiducial simulations; from the figure, we find that removal of such terms has very little effect, implying that leakage into the disconnected

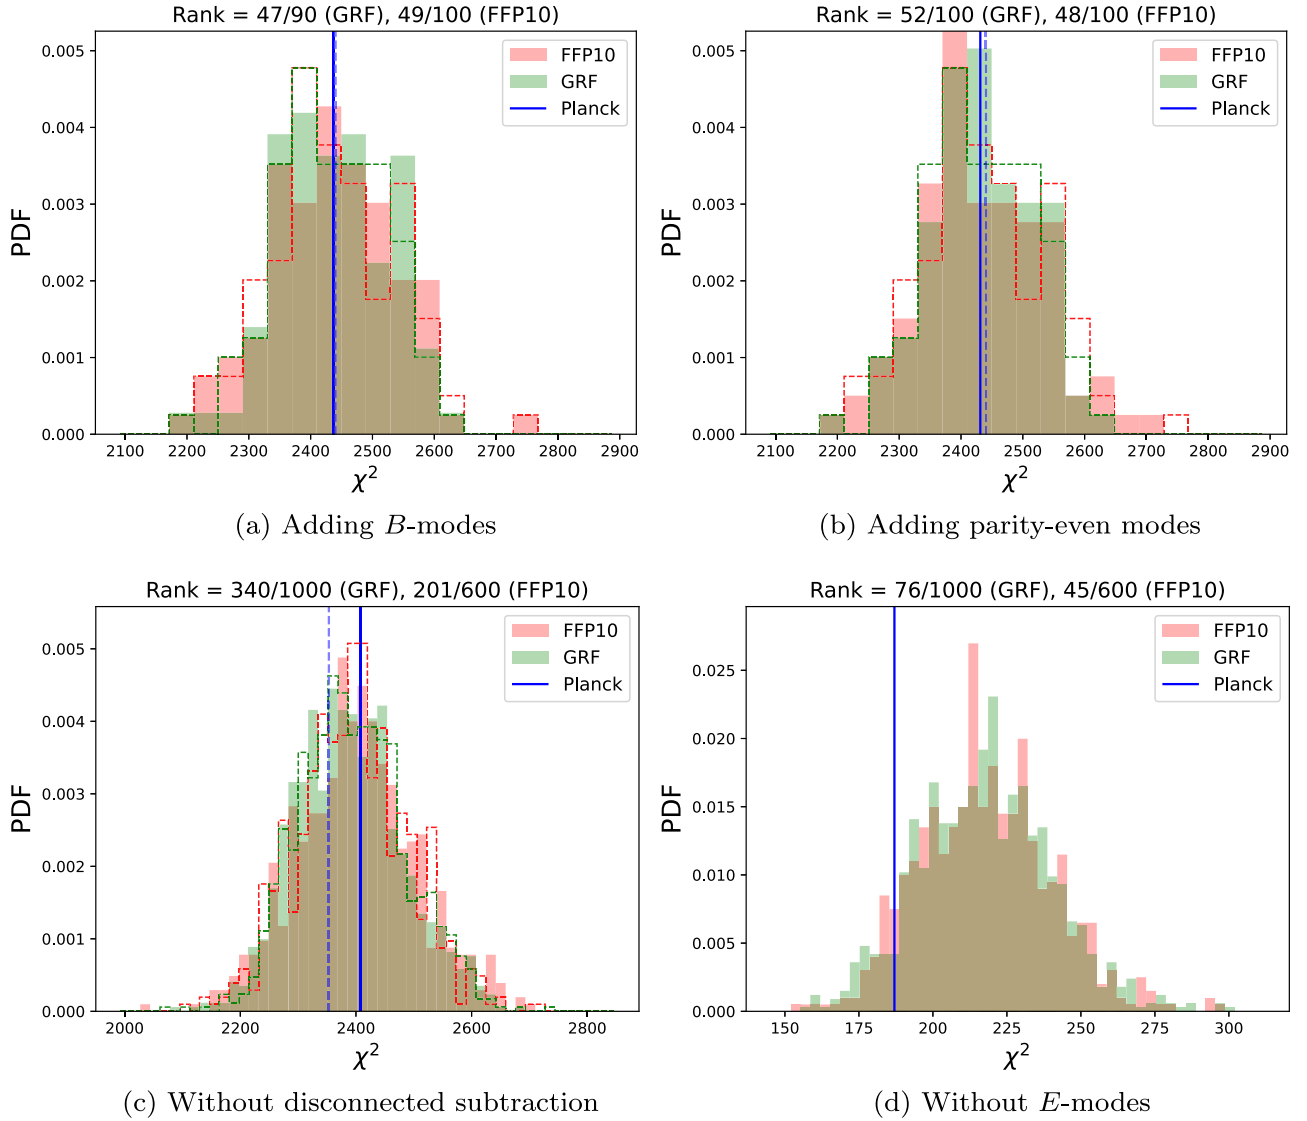


FIG. 10. As Fig. 5, but assessing the consistency of our analyses with various modifications to our fiducial assumptions. In the first three cases (described in Appendix B), we show the χ^2 distributions from the analysis alongside the fiducial results in dashed lines; in the fourth case, we show only the new results, since the dimensionality differs. The first two tests use 100 FFP10 and GRF simulations, whilst the second two use the full fiducial sample. In all cases, the blind tests show no evidence for parity violation.

contributions is small in this case (though more relevant for the gauge-field results of Sec. V). The corresponding detection significance (relative to FFP10) remains at the fiducial value of -0.4σ .

4. Removing polarization

Finally, we consider the effects of removing E -modes from our analysis. This emulates the results of [37], except with slightly broader binning, updated *Planck* PR4

measurements, and including the T - E correlations in the measurement pipeline (by restricting to T modes only after the trispectrum is normalized). In this case, the PDF [shown in Fig. 10(d)], is shifted to much lower χ^2 values, since we reduce the dimensionality from 2386 to 217 bins. Once again, we find no evidence for parity violation, with the *Planck* data consistent with the FFP10 simulations at -1.43σ in a rank test. This is expected, and shows that our new pipeline yields analogous results to the former.

- [1] D. N. Spergel *et al.* (WMAP Collaboration), *Astrophys. J. Suppl. Ser.* **148**, 175 (2003).
- [2] E. Komatsu *et al.* (WMAP Collaboration), *Astrophys. J. Suppl. Ser.* **192**, 18 (2011).
- [3] Planck Collaboration, N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo *et al.*, *Astron. Astrophys.* **641**, A6 (2020).
- [4] Y. Akrami *et al.* (Planck Collaboration), *Astron. Astrophys.* **641**, A9 (2020).
- [5] Y. Akrami *et al.* (Planck Collaboration), *Astron. Astrophys.* **641**, A10 (2020).
- [6] S. Aiola *et al.* (ACT Collaboration), *J. Cosmol. Astropart. Phys.* **12** (2020) 047.
- [7] L. Balkenhol *et al.* (SPT-3G Collaboration), *Phys. Rev. D* **108**, 023510 (2023).
- [8] P. A. R. Ade *et al.* (BICEP/Keck Collaboration), in *56th Rencontres de Moriond on Cosmology* (2022), [arXiv:2203.16556](https://arxiv.org/abs/2203.16556).
- [9] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson, *Phys. Rev.* **105**, 1413 (1957).
- [10] C. Creque-Sarbinowski, S. Alexander, M. Kamionkowski, and O. Philcox, *J. Cosmol. Astropart. Phys.* **11** (2023) 029.
- [11] T. D. Lee and C. N. Yang, *Phys. Rev.* **104**, 254 (1956).
- [12] M. Shiraishi, *Phys. Rev. D* **94**, 083503 (2016).
- [13] S. H. S. Alexander, *Phys. Lett. B* **660**, 444 (2008).
- [14] V. Gluscevic and M. Kamionkowski, *Phys. Rev. D* **81**, 123529 (2010).
- [15] M. Shiraishi, A. Ricciardone, and S. Saga, *J. Cosmol. Astropart. Phys.* **11** (2013) 051.
- [16] N. Bartolo, S. Matarrese, M. Peloso, and M. Shiraishi, *J. Cosmol. Astropart. Phys.* **01** (2015) 027.
- [17] M. Shiraishi, *J. Cosmol. Astropart. Phys.* **06** (2012) 015.
- [18] M. Shiraishi, D. Nitta, and S. Yokoyama, *Prog. Theor. Phys.* **126**, 937 (2011).
- [19] A. Lue, L.-M. Wang, and M. Kamionkowski, *Phys. Rev. Lett.* **83**, 1506 (1999).
- [20] G. Cabass, S. Jazayeri, E. Pajer, and D. Stefanyszyn, *J. High Energy Phys.* **02** (2023) 021.
- [21] G. Cabass, M. M. Ivanov, and O. H. E. Philcox, *Phys. Rev. D* **107**, 023523 (2023).
- [22] S. Alexander and N. Yunes, *Phys. Rep.* **480**, 1 (2009).
- [23] S. H. S. Alexander, *Int. J. Mod. Phys. D* **25**, 1640013 (2016).
- [24] A. D. Sakharov, *Pis'ma Zh. Eksp. Teor. Fiz.* **5**, 32 (1967).
- [25] T. Liu, X. Tong, Y. Wang, and Z.-Z. Xianyu, *J. High Energy Phys.* **04** (2020) 189.
- [26] H. Abedi, M. Ahmadvand, and S. S. Gousheh, *Phys. Lett. B* **786**, 35 (2018).
- [27] R. N. Cahn, Z. Slepian, and J. Hou, *Phys. Rev. Lett.* **130**, 201002 (2023).
- [28] W. R. Coulton, O. H. E. Philcox, and F. Villaescusa-Navarro, *Phys. Rev. D* **109**, 023531 (2024).
- [29] M. Shiraishi, M. Liguori, and J. R. Fergusson, *J. Cosmol. Astropart. Phys.* **01** (2015) 007.
- [30] P. A. R. Ade *et al.* (Planck Collaboration), *Astron. Astrophys.* **594**, A17 (2016).
- [31] O. H. E. Philcox and M. Shiraishi, *Phys. Rev. D* **109**, 063522 (2024).
- [32] O. H. E. Philcox, *Phys. Rev. D* **106**, 063501 (2022).
- [33] J. Hou, Z. Slepian, and R. N. Cahn, *Mon. Not. R. Astron. Soc.* **522**, 5701 (2023).
- [34] J. R. Eskilt and E. Komatsu, *Phys. Rev. D* **106**, 063503 (2022).
- [35] P. Diego-Palazuelos *et al.*, *Phys. Rev. Lett.* **128**, 091302 (2022).
- [36] J. R. Eskilt, *Astron. Astrophys.* **662**, A10 (2022).
- [37] O. H. E. Philcox, *Phys. Rev. Lett.* **131**, 181001 (2023).
- [38] Y. Akrami *et al.* (Planck Collaboration), *Astron. Astrophys.* **643**, A42 (2020).
- [39] M. Tristram *et al.*, *Astron. Astrophys.* **647**, A128 (2021).
- [40] E. Rosenberg, S. Gratton, and G. Efstathiou, *Mon. Not. R. Astron. Soc.* **517**, 4620 (2022).
- [41] O. H. E. Philcox, *Phys. Rev. D* **108**, 063506 (2023).
- [42] O. H. E. Philcox, *Phys. Rev. D* **107**, 123516 (2023).
- [43] D. M. Regan, E. P. S. Shellard, and J. R. Fergusson, *Phys. Rev. D* **82**, 023520 (2010).
- [44] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelman, *Astrophys. J.* **622**, 759 (2005).
- [45] H. F. Gruetjen, J. R. Fergusson, M. Liguori, and E. P. S. Shellard, *Phys. Rev. D* **95**, 043532 (2017).
- [46] Y. Akrami *et al.* (Planck Collaboration), *Astron. Astrophys.* **641**, A4 (2020).
- [47] K. M. Smith, L. Senatore, and M. Zaldarriaga, [arXiv:1502.00635](https://arxiv.org/abs/1502.00635).
- [48] O. H. E. Philcox, PolyBin: Binned polyspectrum estimation on the full sky, Astrophysics Source Code Library, record [ascl:2307.020](https://arxiv.org/abs/2307.020) (2023).
- [49] K. M. Smith and M. Zaldarriaga, *Mon. Not. R. Astron. Soc.* **417**, 2 (2011).
- [50] A. J. S. Hamilton and M. Tegmark, *Mon. Not. R. Astron. Soc.* **312**, 285 (2000).
- [51] A. Kalaja, P. D. Meerburg, G. L. Pimentel, and W. R. Coulton, *J. Cosmol. Astropart. Phys.* **04** (2021) 050.
- [52] A. Girard, *Numer. Math.* **56**, 1 (1989).
- [53] M. Hutchinson, *Commun. Stat. Simul. Comput.* **19**, 433 (1990).
- [54] N. Sailer, E. Castorina, S. Ferraro, and M. White, *J. Cosmol. Astropart. Phys.* **12** (2021) 049.
- [55] github.com/OliverPhilcox/PolyBin.
- [56] R. N. Cahn and Z. Slepian, *J. Phys. A* **56**, 325204 (2023).
- [57] D. A. Varshalovich, A. N. Moskalev, and V. K. Khersonskii, *Quantum Theory of Angular Momentum* (World Scientific, Singapore, 1988), <https://library.open.org/handle/20.500.12657/50493>.