Do the CMB Temperature Fluctuations Conserve Parity?

Oliver H. E. Philcox*
Center for Theoretical Physics, Department of Physics, Columbia University, New York, New York 10027, USA and Simons Foundation, New York, New York 10010, USA

(Received 23 March 2023; revised 16 August 2023; accepted 22 September 2023; published 3 November 2023)

Observations of the cosmic microwave background (CMB) have cemented the notion that the large-scale Universe is both statistically homogeneous and isotropic. But is it invariant also under reflections? To probe this we require parity-sensitive statistics: for scalar observables, the simplest is the trispectrum. We make the first measurements of the parity-odd scalar CMB, focusing on the large-scale ($2 < \ell < 510$) temperature anisotropies measured by *Planck*. This is facilitated by new quasi-maximum-likelihood estimators for binned correlators, which account for mask convolution and leakage between even- and odd-parity components, and achieve ideal variances within ≈20%. We perform a blind test for parity violation by comparing a χ^2 statistic from *Planck* to theoretical expectations, using two suites of simulations to account for the possible likelihood non-Gaussianity and residual foregrounds. We find consistency at the $\approx 0.4\sigma$ level, yielding no evidence for novel early-Universe phenomena. The measured trispectra allow for a wealth of new physics to be constrained; here, we use them to constrain eight primordial models, including ghost inflation, cosmological collider scenarios, and Chern-Simons gauge fields. We find no signatures of new physics, with a maximal detection significance of 2.0σ. Our results also indicate that the recent parity excesses seen in the BOSS galaxy survey are not primordial in origin, given that the CMB dataset contains roughly 250× more primordial modes, and is far easier to interpret, given the linear physics, Gaussian statistics, and accurate mocks. Tighter CMB constraints can be wrought by including smaller scales (though rotational invariance washes out the flat-sky limit) and adding polarization data.

DOI: 10.1103/PhysRevLett.131.181001

Since their inception in the mid-1960s, analyses of the cosmic microwave background (CMB) have revolutionized modern physics, most notably through the associated development of a cosmological model [1,2]. Though early work focused on the CMB frequency spectrum [3], most contemporary experiments consider the spatial fluctuations, both in temperature and polarization [4–6]. Through the measurement of polyspectra, including the power spectrum, bispectrum, and trispectrum, these can be used to constrain the physics of the early and late Universe, placing constraints on a wide variety of models. Upcoming experiments will significantly tighten these bounds, with much of the new information coming from polarization anisotropies at high multipoles [7,8]. Given this push to smaller scales, it is interesting to ask whether this is sufficient for all models of interest; equivalently, have we exhausted all the information in the large-scale temperature fluctuations probed by full-sky experiments such as *Planck*?

Cosmological parity symmetry provides an intriguing counterexample. In particle physics, it has long been known that weak interactions do not preserve parity (hereafter $\mathbb{P}$) symmetry [9]; cosmologically, most interactions of interest are gravitational, and thus $\mathbb{P}$ invariant. In the early Universe, this is not the case, and there are theoretical hints that parity violation could accompany phenomena such as baryogenesis (via some leptogenesis mechanism) [10–13]. If this occurs during inflation (with examples involving massive exchange particles [14] and modified gravity [10]), signatures would naturally be left in the distribution of matter and gravitational waves. In the latter case, $\mathbb{P}$-violating power spectra such as TB and EB would be generated [15], or higher-order bispectra such as TTB [20,22–30]. For scalar observables, such as temperature fluctuations, $\mathbb{P}$-violating physics can only be seen in a temperature trispectrum or above, due to the equivalence of parity transformations and three-dimensional rotations for low order statistics (first discussed for galaxy survey contexts in [31]). This has been noted in previous works [32–36], but never explicitly searched for in CMB data [37].

Determining whether the CMB obeys parity symmetry is a topic of great current relevance. Recently, tentative evidence for $\mathbb{P}$ violation has been reported using CMB birefringence [38–40], as well as the large-scale distribution of galaxies [33,41], with the latter using the approach first proposed in [31]. Given that systematic effects such as unaccounted-for cosmic dust [42] or unjustified analysis choices (e.g., a Gaussian likelihood or mocks of insufficient fidelity) could be an alternative explanation for such results, full confirmation requires an independent probe. In this Letter, we perform the first tests of scalar $\mathbb{P}$ violation in the CMB by measuring the four-point correlator of the CMB

temperature anisotropies, using methods similar to those of other primordial non-Gaussianity analyses (see Ref. [43] and references therein), as well as new tools for estimating the correlators themselves [44]. Since this is a scalar observable, it cannot be used to directly constrain birefringence via axion-photon interactions [12,23,45,46]); these occur at late times, and affect only photon polarization. However, it is a direct probe of the primordial density perturbations, and is thus sensitive to potential early Universe processes that could source parity violation in the distribution of galaxies. Indeed, the CMB is likely a stronger probe, given its large signal-to-noise and almost-Gaussian statistics.

Given this novel dataset, we may ask two questions: (i) do we see *any* evidence of parity violation in the large-scale *Planck* temperature map? (ii) what bounds can we place on physical models of parity violation? The latter can be compared to constraints from galaxy surveys [34], with the current results benefiting from a much larger maximum scale ($\ell \approx 3$), the cosmic-variance limit of the CMB, and the well understood statistics and covariances. Below, we present this analysis, with additional details given in the Supplemental Material [47]. To avoid confirmation bias, the entire (public) analysis pipeline (including trispectrum estimation, model computation, and null test pipeline) was developed and tested on realistic simulations before the *Planck* data were analyzed.

Theoretical framework.—The reduced CMB trispectrum, $t^{\ell_1\ell_2}_{\ell_3\ell_4}(L)$, is defined as

$$\left\langle \prod_{i=1}^{4} a_{\ell_i m_i} \right\rangle_c \equiv \sum_{LM} (-1)^M w^{L(-M)}_{\ell_1\ell_2 m_1 m_2} w^{LM}_{\ell_3\ell_4 m_3 m_4} t^{\ell_1\ell_2}_{\ell_3\ell_4}(L) + 23 \text{ perms.} \tag{1}$$

[44,48], where $a_{\ell m}$ is the primary CMB temperature fluctuation, defining the weights

$$w^{LM}_{\ell_1\ell_2 m_1 m_2} \equiv \sqrt{\frac{(2\ell_1+1)(2\ell_2+1)(2L+1)}{4\pi}} \times \begin{pmatrix} \ell_1 & \ell_2 & L \\ m_1 & m_2 & -M \end{pmatrix} \begin{pmatrix} \ell_1 & \ell_2 & L \\ -1 & -1 & 2 \end{pmatrix} \tag{2}$$

(choosing spins as in [44] to avoid parity-odd cancellations). Here, the 3×2 matrices are Wigner $3j$ symbols, and the trispectrum is indexed by the multipoles of four sides, ℓ_i, and a diagonal, L, with $\{\ell_1, \ell_2, L\}$ and $\{\ell_3, \ell_4, L\}$ obeying triangle conditions. Trispectra with even (odd) $\ell_1 + \ell_2 + \ell_3 + \ell_4$ probe the parity-even (parity-odd) contributions to the CMB perturbations; we consider the latter in this work.

The CMB temperature fluctuations are a direct probe of the underlying primordial curvature perturbation, $\zeta(\mathbf{k})$:

$$a_{\ell m} \equiv 4\pi i^\ell \int_{\mathbf{k}} \zeta(\mathbf{k}) \mathcal{T}_\ell(k) Y^*_{\ell m}(\hat{\mathbf{k}}), \tag{3}$$

where $Y_{\ell m}$ is a spherical harmonic, $\int_{\mathbf{k}} \equiv \int d^3k/(2\pi)^3$, and $\mathcal{T}_\ell$ is the usual transfer function, encoding radiation physics. It follows that the CMB trispectrum is linearly related to the curvature trispectrum, T_ζ: as such, it can be used to place constraints on various models of inflationary parity violation.

Data analysis.—Our dataset is the full-mission *Planck* 2018 SMICA component-separated temperature map [43,49]. [50] We additionally use a set of 300 FFP10 simulations [51,52] to validate our pipeline; these are generated at a cosmology similar to [53] and processed with the same SMICA pipeline. To model the trispectrum noise properties, we generate a suite of 1000 Gaussian random field (hereafter GRF) maps at a HEALPIX $N_{\rm side} = 2048$ [54] with the *Planck* 2018 best-fit cosmology and noise properties [43]. To begin our analysis, we downgrade all maps to $N_{\rm side} = 256$ (for computational efficiency), then modulated them by a linear filter, denoted S^{-1}, similar to the approach of [55] and [56]. This applies the common sky mask described in [43] (with 10-arcminute smoothing, yielding $f_{\rm sky} = 0.79$) to null the brightest foreground regions and linearly inpaints small holes (those containing fewer than 40 pixels), as described in [57]. Maps are then filtered in harmonic space (at $N_{\rm side} = 256$) using a $C^{TT}_\ell + N_\ell$ transfer function, where C^{TT}_ℓ is the theoretical power spectrum evaluated at the best-fit cosmology of [2], and the noise spectrum N_ℓ is extracted from the SMICA half-mission maps.

Given the data, the S^{-1} weighting scheme, and the *Planck* beam (including the pixel window function), correlators are obtained using the POLYBIN code (see Ref. [44] for discussion and verification) [58]. This implements the binned trispectrum estimator, $\hat{t}$, from a quartic estimator ($\hat{t}^{\rm num}$, with residual disconnected contributions subtracted using 100 GRF simulations) and a data-independent normalization matrix $\mathcal{F}$:

$$\hat{t}(\mathbf{b}, B) = \sum_{\mathbf{b}', B'} \mathcal{F}^{-1}(\mathbf{b}, B; \mathbf{b}', B') \hat{t}^{\rm num}(\mathbf{b}', B'), \tag{4}$$

where $\mathbf{b}$, B are bins in $\{\ell_i\}$ and L, respectively (cf. [60,61] in 3D). $\mathcal{F}$ is estimated using 50 Monte Carlo realizations (using tricks analogous to [56]), and found to be extremely well converged [44]. Explicit definitions of each component can be found in [44]. This is a quasi-maximum-likelihood estimator, which, through the normalization matrix (which is approximately the variance of $\hat{t}^{\rm num}$), accounts for leakage between bins and components of different parity, being unbiased for any choice of S^{-1}. [62] The output spectra are close to minimum variance (with equality obtained in the limit of Gaussian maps and S^{-1} equal to the true inverse pixel covariance), can be

(approximately) compared to theory without mask convolution (cf. [64] for the power spectrum), and can be efficiently computed [with $\mathcal{O}(N_{\rm bins})$ complexity]. In limiting regimes, the inverse of the normalization matrix is equal to the analytic trispectrum covariance (i.e., it is a Fisher matrix); this motivates working with the transformed variable

$$\hat{\tau}(\mathbf{b}, B) \equiv \sum_{\mathbf{b}', B'} \mathcal{F}^{\mathrm{T}/2}(\mathbf{b}, B; \mathbf{b}', B') \hat{t}(\mathbf{b}', B') \tag{5}$$

(for Cholesky factorization $\mathcal{F}^{1/2}$ and transpose T), which is close to a unit normal variable, cf. [65]. Since $\mathcal{F}$ well describes the correlation structure of the data (validated in the Supplemental Material [47]), this accounts for correlations between bins (induced by the mask and cosmic variance), such that the individual τ bins are uncorrelated.

Noting that the trispectrum dimensionality grows as N_ℓ^5 (for N_ℓ bins per dimension), and the expected squared signal-to-noise scales as $\ell_{\rm max}^2$ (in the scale-invariant limit, cf. [66]), we adopt a uniform binning in $\ell^{2/5}$, including 9 bins in the range [2, 510] [67], dropping any configurations whose bin centers do not satisfy the triangle conditions. We fix $\ell_{\rm min} = 2$ to avoid vanishing weights in (2). Our choice of $\ell_{\rm max}$ (and thus $N_{\rm side}$) is motivated both by computational considerations, [70] and noting that parity-odd spectra are heavily suppressed at high-ℓ due to projection effects and Silk damping [66]. To minimize leakage from unmeasured bins and foreground contamination, we simultaneously measure the numerators of both parity-even and parity-odd trispectra (a total of 1273 components, given the above bin restrictions), then discard the former (and the first bin) after the deconvolving normalization matrix has been applied (4). This leaving 460 parity-odd bins with $\ell_{\rm min} = 3$, $\ell_{\rm max} = 510$. Computation of the (data-independent) normalization matrix $\mathcal{F}$ requires ≈1000 CPU-h (and 200 GB memory, without optimization), with each simulation requiring an additional 10 CPU-h to analyze. In the Supplemental Material [47], we show the normalization matrices, $\mathcal{F}$, and present Fisher forecasts examining the impact of $\ell_{\rm max}$ on parameter constraints.

Blind tests for parity violation.—To examine whether the *Planck* data contains evidence for $\mathbb{P}$ violation, we first perform a blind model-agnostic test. Here, we follow a methodology similar to [31,33,41,71], first computing the normalized trispectrum $\tau(\mathbf{b}, B)$ via (5). This is shown in Fig. 1. For both the FFP10 and GRF simulations, we find a variance relatively close to unity, indicating that our estimators are close to optimal (primarily within ≈20%, though with slightly larger deviations on small scales). Deviations arise from (a) insufficient simulations used to remove disconnected components, (b) suboptimal weights S^{-1} (which could be rectified by a conjugate-gradient-inversion scheme), (c) insufficient GRF simulations used to estimate the normalization matrix, and (d) likelihood non-Gaussianity (but not large-scale foregrounds, as evidenced by the same behavior for the two sets of simulations). We have verified that (a) can be important (with larger variances seen when halving the number of bias simulations), but (c) is negligible, with a well-converged normalization matrix found even with five Monte Carlo realizations. Notably, suboptimal variances do not bias our analyses, cf. [44], though slightly degrade their constraining power. In the Supplemental Material [47], we verify that the correlation matrix does not show any noticeable departures from the identity.

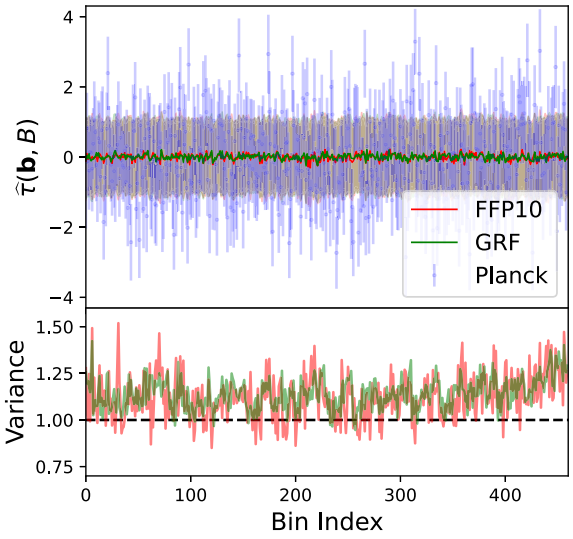

FIG. 1. Normalized parity-odd trispectrum, $\tau(\mathbf{b}, B)$ measured from *Planck* data (blue), 300 FFP10 simulations (red), and 1000 GRF simulations (green) using the PolyBin code with $\ell_{\rm min} = 3$, $\ell_{\rm max} = 510$. The horizontal axis gives the bin index (with large-scale modes on the left, ordered with $\ell_1 \le \ell_2, \ell_1 \le \ell_3, \ell_3 \le \ell_4$), and shaded regions show the 1σ variance from simulations (shown explicitly in the bottom panel). The variance is close to the Gaussian prediction (unity) but exhibits a slight enhancement due to the nontrivial mask and finite number of bias simulations, though the correlation matrix is consistent with the identity (see the Supplemental Material [47]).

Data are analyzed via the following χ^2 statistic:

$$\hat{\chi}^2 = \sum_{\mathbf{b}, B} \frac{\hat{\tau}^2(\mathbf{b}, B)}{\mathrm{var}[\tau(\mathbf{b}, B)]}, \tag{6}$$

where the variance is computed from the GRF simulations, noting that the individual bins of τ are uncorrelated (which obviates the need for additional data compression). Under the following assumptions, the statistic can be analyzed using a χ^2 distribution with 460 degrees of freedom: (i) t obeys Gaussian statistics; (ii) the GRF simulations accurately reproduce the *Planck* trispectrum covariance. To test (i), we implement rank tests, comparing $\hat{\chi}^2$ to the empirical χ^2 distribution from the GRF (or FFP10) simulations. For (ii), we note that the GRF simulations adopt a cosmology

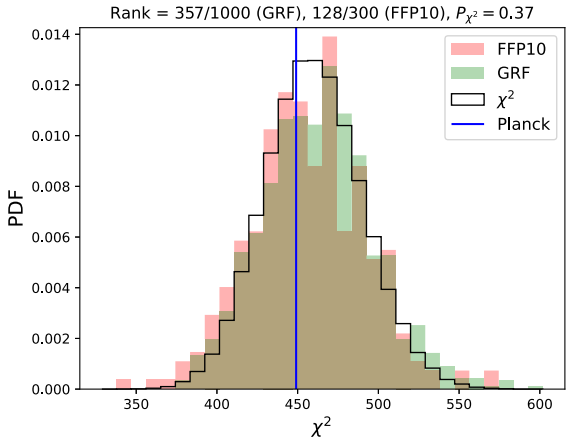

FIG. 2. Probability of parity violation in the *Planck* dataset. We plot the χ^2 values (6) extracted from the *Planck* SMICA maps (blue), the FFP10 simulations (red), and GRF simulations (green), alongside the binned theoretical prediction (black). Deviation of the *Planck* results from the expected curves would give evidence for parity violation. Relative to the GRF simulations, the *Planck* data has a rank of $357/1000$ (i.e., a larger χ^2 than that of 356 simulations), or $128/300$ with respect to FFP10, with the consistency of the two sets of simulations indicating that our pipeline is robust and not affected by residual foregrounds. For the theoretical distribution, we find a χ^2 probability to exceed of 63%, though caution that the likelihood may be non-Gaussian. Overall, we find no evidence for parity violation in the *Planck* dataset.

close to *Planck*, and, unlike for LSS, higher-point functions entering the trispectrum covariance are negligible on the scales of interest. To test this assumption, we look for differences between the FFP10 and GRF results; we find excellent agreement, indicating that deviations of $\mathrm{var}(\tau)$ from unity arise primarily from masking and the finite number of bias realizations, which affect the simulations and data identically, i.e., there is no evidence for contamination from residual foregrounds. If the *Planck* $\hat{\chi}^2$ value lies significantly north of its expectation, we can conclude that there is evidence for large-scale $\mathbb{P}$ violation, since no conventional CMB physics would generate such a parity-odd signature. In contrast, a low $\hat{\chi}^2$ value would imply systematic contamination.

Figure 2 shows the χ^2 values extracted from the *Planck* data, the 300 FFP10 simulations and the 1000 GRF realizations. Before discussing the data, we first note that the simulated distributions exhibit slightly wider posteriors than the analytic prediction; this implies that the trispectrum likelihood is mildly non-Gaussian [72]. If not accounted for (particularly in the distribution tails), this could lead to a false detection of parity violation. Second, we find similar empirical distributions from the simulation suites: relative to the GRF simulations, the FFP10 distribution has a mean probability to exceed (PTE) of $(46 \pm 2)\%$, relative to the expectation $(50 \pm 2)\%$, implying that any residual foreground-induced bias is small (though we note that the simulations do not include all correlations between foregrounds, which are relevant at large ℓ, cf. [73]).

We now consider the *Planck* data. From the raw measurements shown in Fig. 1 it is difficult to draw firm conclusions; however, Fig. 2 allows us to compare the observed χ^2 to the simulated distributions. Relative to the GRF suite, we find a rank of $357/1000$ (equivalent to a PTE of 64% or -0.4σ) or $128/300$ (equivalent to $p = 57\%$ or -0.2σ) with respect to FFP10. The conclusions are clear: the large-scale *Planck* temperature trispectrum shows no evidence for parity violation.

Discussion.—In this Letter, we have presented constraints on parity-violating physics from the large-scale *Planck* temperature, through a blind test for $\mathbb{P}$ violation and (in the Supplemental Material [47], which includes [74–83]) constraints on the amplitudes of various physical models. We now consider the physical implications of the above.

Previous constraints on scalar parity violation were obtained from LSS in [33,41]. To compare the two sets of results, it is useful to consider the number of linear modes present in each analysis. For the CMB, this is straightforward:

$$N_{\mathrm{modes}}^{\mathrm{CMB}} = f_{\mathrm{sky}} \sum_{\ell=\ell_{\mathrm{min}}}^{\ell_{\mathrm{max}}} (2\ell+1) \left[\frac{C_\ell^{\mathrm{TT}}}{C_\ell^{\mathrm{TT}} + N_\ell^{\mathrm{TT}}}\right]^2, \quad (7)$$

giving $\approx 2 \times 10^5$ modes for $\ell_{\mathrm{max}} = 510$ (with the signal-to-noise scaling roughly supported by the Fisher forecasts described in the Supplemental Material [47], implying that projection effects are subdominant on these scales). For galaxy surveys, we must account for the decorrelation of modes due to gravitational evolution. Using the approach of [84], we can estimate the number of remaining linear modes at a fixed redshift z as

$$N_{\mathrm{modes}}^{\mathrm{LSS}} = V \int_{k_{\mathrm{min}}}^{k_{\mathrm{max}}} \frac{k^2 dk}{2\pi^2} \frac{d\mu}{2} \left[\frac{G^2(k,\mu,z) P_{\mathrm{lin}}(k,\mu,z)}{P_{\mathrm{tot}}(k,\mu,z)}\right]^2, \quad (8)$$

where V is the survey volume, P_{lin} and P_{tot} are the linear theory and total power spectra, and $G(k,\mu,z) \equiv \langle \delta_{\mathrm{lin}} \delta_{\mathrm{tot}} \rangle / \langle \delta_{\mathrm{lin}} \delta_{\mathrm{lin}} \rangle$ gives the decorrelation of modes from the initial conditions, estimated via a Zel'dovich propagator. Setting $k \in 2\pi/[160, 20] h\,\mathrm{Mpc}^{-1}$ and $V \approx 3.9 h^{-3}\,\mathrm{Gpc}^3$, matching [33,41], we find a total of $\approx 10^3$ linear modes, or $\approx 10^5$ if we ignore the propagator term, and assume all modes trace the initial conditions.

The above comparison indicates that the CMB dataset is a much stronger probe of the parity-violating initial conditions. As such, our blind-test results suggest that the tentative detection of $\mathbb{P}$ violation seen in [33,41] is not primordial (since the 2.9σ hint of [33] would correspond to

around 50σ in the CMB, based on primordial scaling arguments alone): instead, it could arise from exotic late-time physics (though see the caveats in [34]), experimental artifacts, or analysis systematics (such as mocks of insufficient quality or likelihood non-Gaussianity).

Turning to the constraints on physical models described in the Supplemental Material [47], we have found no evidence for any specific inflationary parity-breaking paradigm, with a maximal detection significance of 2.0σ. Our limits on the gauge model are broadly consistent with those forecasted in [32] (at the level appropriate for the inexact nature of such predictions). Furthermore, the Fisher forecasts performed in the Supplemental Material [47] yield similar constraints to the full analysis, implying that our posteriors are desirably Gaussian. Comparing to constraints from LSS, we find somewhat weaker bounds than those of former works [33,34]. This may appear to lie in contrast with the above mode-counting hierarchy; however, the LSS constraints did not account for mode decorrelation (instead assuming all modes were in the linear regime) or nonlinear structure growth (which alters the scale dependence of the signature), and did not marginalize over galaxy formation uncertainties. Each effect can lead to an inflation of the error bars by several orders of magnitude [85]. Furthermore, the precise constraints on physical models should be considered with caution since numerical inaccuracies in the template computation can lead to falsely enhanced precision (since noise in τ always increases the normalization matrix)—the numerical convergence was more carefully considered for the templates used herein than for the LSS case.

Though the primordial templates generated by the three classes of models considered herein span a wide variety of the primordial space, they are not all encompassing. Alternative models could include strong breaking of scale-invariance [87], the exchange of higher spin particles, or the inclusion of inflationary modified gravity (e.g., with Chern-Simons interactions [10,88,89]). The bounds on such models can be straightforwardly computed given the relevant CMB templates, and proceeds analogously to the above analysis. An alternative option would be to devise a general basis for primordial parity-odd trispectra onto which any model can be projected, in analogy to the $f_{\rm NL}$ parametrization of bispectra.

Finally, we consider future prospects. On large scales, the primary CMB is cosmic-variance limited, such that the signal-to-noise on inflationary models will not increase with future experiments. However, one can consider the addition of polarization data (in particular E modes), which can add substantial information on scalar physics [90]; in the case of $f_{\rm NL}^{\rm eq}$, for example, polarization provides $\approx 60\%$ of the constraining power [43]. This will require modifying the trispectrum estimators of [44] to include fields of nonzero spin, but is not conceptually more complicated [91]. Furthermore, one can extend the current analysis to smaller scales (for example, making use of high-resolution ACT and SPT data), albeit at the cost of increased dimensionality and computation time. From the Fisher forecasts performed in the Supplemental Material [47], doubling the ℓ range could increase the squared signal to noise by a factor of 4 (depending on the model in question), though we caution that projection effects start to limit analyses by $\ell_{\rm max} \sim 10^3$ [66]. All in all, while these are the first constraints on CMB scalar parity violation, they seem unlikely to be the last.

We thank Stephon Alexander, William Coulton, Cyril Creque-Sarbinowski, Adriaan Duivenvoorden, Giulio Fabian, Colin Hill, Marc Kamionkowski, Daniel Meerburg, and William Underwood for insightful discussions that led to this work, as well as Neal Dalal, Krzysztof Gorski, Meng-Xiang Lin, Minh Nguyen, Ue-Li Pen, and Maresuke Shiraishi for postsubmission feedback. We are additionally grateful to William Coulton, Cyril Creque-Sarbinowski, Colin Hill, and David Spergel for careful reading of the manuscript, as well as the anonymous referees for insightful feedback. O. H. E. P. thanks the Simons Society of Fellows and thanks John Watling of Nassau for support. The author is pleased to acknowledge that the work reported in this Letter was substantially performed using the Princeton Research Computing resources at Princeton University, which is a consortium of groups led by the Princeton Institute for Computational Science and Engineering (PICSciE) and the Office of Information Technology's Research Computing Division.

*ohep2@cantab.ac.uk

[1] D. N. Spergel *et al.* (WMAP Collaboration), Astrophys. J. Suppl. Ser. **148**, 175 (2003).
[2] N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo *et al.* (Planck Collaboration), Astron. Astrophys. **641**, A6 (2020).
[3] D. J. Fixsen, E. S. Cheng, J. M. Gales, J. C. Mather, R. A. Shafer, and E. L. Wright, Astrophys. J. **473**, 576 (1996).
[4] S. Aiola *et al.* (ACT Collaboration), J. Cosmol. Astropart. Phys. 12 (2020) 047.
[5] P. A. R. Ade *et al.* (BICEP2, Keck Array Collaborations), Phys. Rev. Lett. **116**, 031302 (2016).
[6] B. A. Benson *et al.* (SPT-3G Collaboration), Proc. SPIE Int. Soc. Opt. Eng. **9153**, 91531P (2014).
[7] P. Ade *et al.* (Simons Observatory Collaboration), J. Cosmol. Astropart. Phys. 02 (2019) 056.
[8] K. N. Abazajian *et al.* (CMB-S4 Collaboration), arXiv:1610.02743.
[9] T. D. Lee and C. N. Yang, Phys. Rev. **104**, 254 (1956).
[10] S. Alexander, A. Marciano, and D. Spergel, J. Cosmol. Astropart. Phys. 04 (2013) 046.
[11] S. Alexander and P. Meszaros, arXiv:hep-th/0703070.
[12] S. H. S. Alexander, Int. J. Mod. Phys. D **25**, 1640013 (2016).

[13] A. D. Sakharov, Pis'ma Zh. Eksp. Teor. Fiz. **5**, 32 (1967).
[14] N. Arkani-Hamed and J. Maldacena, arXiv:1503.08043.
[15] These can also be generated by anisotropic (parity-conserving) inflationary models, though this generates only off-diagonal contributions [16–20]. Gravitational wave effects can also generate scalar four-point functions via "fossil" effects, cf. [21], though these are constrained from their CMB *B*-mode signatures.
[16] M. Watanabe, S. Kanno, and J. Soda, Phys. Rev. Lett. **102**, 191302 (2009).
[17] M. Watanabe, S. Kanno, and J. Soda, Mon. Not. R. Astron. Soc. **412**, L83 (2011).
[18] A. Nicolis and G. Sun, J. Cosmol. Astropart. Phys. 04 (2021) 074.
[19] A. Emir Gümrükçüoğlu, B. Himmetoglu, and M. Peloso, Phys. Rev. D **81**, 063528 (2010).
[20] N. Bartolo, S. Matarrese, M. Peloso, and M. Shiraishi, J. Cosmol. Astropart. Phys. 01 (2015) 027.
[21] K. W. Masui, U.-L. Pen, and N. Turok, Phys. Rev. Lett.**118**, 221301 (2017).
[22] A. Lue, L.-M. Wang, and M. Kamionkowski, Phys. Rev. Lett. **83**, 1506 (1999).
[23] V. Gluscevic and M. Kamionkowski, Phys. Rev. D **81**, 123529 (2010).
[24] M. Kamionkowski and T. Souradeep, Phys. Rev. D **83**, 027301 (2011).
[25] M. Shiraishi, A. Ricciardone, and S. Saga, J. Cosmol. Astropart. Phys. 11 (2013) 051.
[26] M. Shiraishi, J. Cosmol. Astropart. Phys. 11 (2013) 006.
[27] M. Shiraishi, J. Cosmol. Astropart. Phys. 06 (2012) 015.
[28] M. Shiraishi, D. Nitta, and S. Yokoyama, Prog. Theor. Phys. **126**, 937 (2011).
[29] N. Bartolo, G. Orlando, and M. Shiraishi, J. Cosmol. Astropart. Phys. 01 (2019) 050.
[30] D. Molinari, Proc. Sci. FRAPWS2018 (**2018**) 021.
[31] R. N. Cahn, Z. Slepian, and J. Hou, Phys. Rev. Lett. **130**, 201002 (2023).
[32] M. Shiraishi, Phys. Rev. D **94**, 083503 (2016).
[33] O. H. E. Philcox, Phys. Rev. D **106**, 063501 (2022).
[34] G. Cabass, M. M. Ivanov, and O. H. E. Philcox, Phys. Rev. D **107**, 023523 (2023).
[35] S. H. S. Alexander, Phys. Lett. B **660**, 444 (2008).
[36] W. R. Coulton, O. H. E. Philcox, and F. Villaescusa-Navarro, arXiv:2306.11782.
[37] Note that new physics models generically modify the scalar and tensor CMB in different ways, thus it remains useful to look at the temperature four-point function after a non-detection of the *B*-mode two-point functions.
[38] P. Diego-Palazuelos *et al.*, Phys. Rev. Lett. **128**, 091302 (2022).
[39] J. R. Eskilt, Astron. Astrophys. **662**, A10 (2022).
[40] J. R. Eskilt and E. Komatsu, Phys. Rev. D **106**, 063503 (2022).
[41] J. Hou, Z. Slepian, and R. N. Cahn, Mon. Not. R. Astron. Soc. **522**, 5701 (2023).
[42] S. E. Clark, C.-G. Kim, J. C. Hill, and B. S. Hensley, Astrophys. J. **919**, 53 (2021).
[43] Y. Akrami *et al.* (Planck Collaboration), Astron. Astrophys. **641**, A9 (2020).
[44] O. H. E. Philcox, Phys. Rev. D **107**, 123516 (2023).
[45] S. M. Carroll, Phys. Rev. Lett. **81**, 3067 (1998).
[46] M. S. Turner and L. M. Widrow, Phys. Rev. D **37**, 2743 (1988).
[47] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevLett.131.181001 for constraints on parity-breaking models, derivations, and other associated analysis material.
[48] D. M. Regan, E. P. S. Shellard, and J. R. Fergusson, Phys. Rev. D **82**, 023520 (2010).
[49] Y. Akrami *et al.* (Planck Collaboration), Astron. Astrophys. **641**, A4 (2020).
[50] Available at http://pla.esac.esa.int/.
[51] R. Adam *et al.* (Planck Collaboration), Astron. Astrophys. **594**, A9 (2016).
[52] N. Aghanim *et al.* (Planck Collaboration), Astron. Astrophys. **641**, A3 (2020).
[53] P. A. R. Ade *et al.* (Planck Collaboration), Astron. Astrophys. **594**, A13 (2016).
[54] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelman, Astrophys. J. **622**, 759 (2005).
[55] P. A. R. Ade *et al.* (Planck Collaboration), Astron. Astrophys. **594**, A17 (2016).
[56] K. M. Smith, L. Senatore, and M. Zaldarriaga, arXiv:1502.00635.
[57] H. F. Gruetjen, J. R. Fergusson, M. Liguori, and E. P. S. Shellard, Phys. Rev. D **95**, 043532 (2017).
[58] Available at https://github.com/oliverphilcox/PolyBin [59].
[59] O. H. E. Philcox, PolyBin: Binned polyspectrum estimation on the full sky, https://ui.adsabs.harvard.edu/abs/2023ascl.soft07020P
[60] O. H. E. Philcox, Phys. Rev. D **103**, 103504 (2021).
[61] O. H. E. Philcox, Phys. Rev. D **104**, 123529 (2021).
[62] Note that we assume the data and mask to be uncorrelated (since the signal is from recombination); see [63] for discussion of when this assumption breaks down and its potential amelioration.
[63] K. M. Surrao, O. H. E. Philcox, and J. Colin Hill, arXiv:2302.05436.
[64] E. Hivon, K. M. Gorski, C. B. Netterfield, B. P. Crill, S. Prunet, and F. Hansen, Astrophys. J. **567**, 2 (2002).
[65] A. J. S. Hamilton, Lect. Notes Phys. **665**, 433 (2008).
[66] A. Kalaja, P. D. Meerburg, G. L. Pimentel, and W. R. Coulton, J. Cosmol. Astropart. Phys. 04 (2021) 050.
[67] The contiguous bins are defined by the edges {2,3,13,35,71,121,189,275,382,510}. Notably, the signal-to-noise scaling depends on the model in question as discussed in [66] and the Supplemental Material [47] (including [68,69]), thus our binning may be suboptimal. For an analysis centered on a particular model type (e.g., a contact trispectrum), one would want to choose the binning accordingly, e.g., dropping all but collapsed tetrahedra.
[68] N. Kogo and E. Komatsu, Phys. Rev. D **73**, 083007 (2006).
[69] L. Bordin and G. Cabass, J. Cosmol. Astropart. Phys. 06 (2019) 050.
[70] Theoretical computation of the trispectrum templates scales at least as $\ell_{\rm max}^3$, while, at fixed number of bins, the trispectrum computation scales with $N_{\rm side}^2$ (for each harmonic transform) and thus $\ell_{\rm max}^2$. Restricting to particular trispectrum configurations (e.g., collapsed or squeezed) can

enable larger $\ell_{\rm max}$ without prohibitive computation times, though this goes against the spirit of a blind null test.
[71] O. H. E. Philcox, J. Hou, and Z. Slepian, arXiv:2108.01670.
[72] In initial testing using $\ell_{\rm max} = 245$, but the same number of bins, this effect was significantly amplified.
[73] J. C. Hill, Phys. Rev. D **98**, 083542 (2018).
[74] N. Arkani-Hamed, P. Creminelli, S. Mukohyama, and M. Zaldarriaga, J. Cosmol. Astropart. Phys. 04 (2004) 001.
[75] M. Bucher, B. Racine, and B. van Tent, J. Cosmol. Astropart. Phys. 05 (2016) 055.
[76] P. A. R. Ade, N. Aghanim, C. Armitage-Caplan, M. Arnaud, M. Ashdown, F. Atrio-Barandela, J. Aumont, C. Baccigalupi, A. J. Banday *et al.* (Planck Collaboration), Astron. Astrophys. **571**, A24 (2014).
[77] D. Alonso, J. Sanchez, and A. Slosar (LSST Dark Energy Science Collaboration), Mon. Not. R. Astron. Soc. **484**, 4127 (2019).
[78] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, Publ. Astron. Soc. Pac. **125**, 306 (2013).
[79] A. E. Bayer and U. Seljak, J. Cosmol. Astropart. Phys. 10 (2020) 009.
[80] R. N. Cahn and Z. Slepian, J. Phys. A **56**, 325204 (2023).
[81] D. A. Varshalovich, A. N. Moskalev, and V. K. Khersonskii, *Quantum Theory of Angular Momentum* (World Scientific, Singapore, 1988).
[82] R. Mehrem, J. T. Londergan, and M. H. Macfarlane, J. Phys. A **24**, 1435 (1991).
[83] X. Niu, M. H. Rahat, K. Srinivasan, and W. Xue, J. Cosmol. Astropart. Phys. 05 (2023) 018.
[84] N. Sailer, E. Castorina, S. Ferraro, and M. White, J. Cosmol. Astropart. Phys. 12 (2021) 049.
[85] The latter is demonstrated in [86], whereupon $f_{\rm NL}$ constraints are inflated by $\sim 100\times$ due to nuisance parameter marginalization even for the optimistic case of a high-redshift sample analyzed jointly with power spectrum data.
[86] G. Cabass, M. M. Ivanov, O. H. E. Philcox, M. Simonovic, M. Zaldarriaga, arXiv:2211.14899.
[87] G. Cabass, S. Jazayeri, E. Pajer, and D. Stefanyszyn, J. High Energy Phys. 02 (2023) 021.
[88] S. Alexander and N. Yunes, Phys. Rep. **480**, 1 (2009).
[89] C. Creque-Sarbinowski, S. Alexander, M. Kamionkowski, and O. Philcox, arXiv:2303.04815.
[90] O. H. E. Philcox and M. Shiraishi, arXiv:2308.03831.
[91] O. H. E. Philcox, Phys. Rev. D **108**, 063506 (2023).