# Probing parity violation with the four-point correlation function of BOSS galaxies

Oliver H. E. Philcox*

*Department of Astrophysical Sciences, Princeton University, Princeton, New Jersey 08540, USA and School of Natural Sciences, Institute for Advanced Study, 1 Einstein Drive, Princeton, New Jersey 08540, USA*

(Received 8 June 2022; accepted 26 July 2022; published 6 September 2022; corrected 16 February 2023)

Parity-violating physics in the early universe can leave detectable traces in late-time observables. While vector- and tensor-type parity violation can be observed in the $B$-modes of the cosmic microwave background, scalar-type signatures are visible only in the four-point correlation function (4PCF) and beyond. This work presents a blind test for parity violation in the 4PCF of the baryon oscillation spectroscopic survey (BOSS) CMASS sample, considering galaxy separations in the range $[20, 160]\ h^{-1}\,\mathrm{Mpc}$. The parity-odd 4PCF contains no contributions from standard ΛCDM physics and can be efficiently measured using recently developed estimators. Data are analyzed using both a nonparametric rank test (comparing the BOSS 4PCFs to those of realistic simulations) and a compressed $\chi^2$ analysis, with the former avoiding the assumption of a Gaussian likelihood. These find similar results, with the rank test giving a detection probability of 99.6% ($2.9\sigma$). This provides significant evidence for parity violation, from either cosmological sources or systematics. We perform a number of systematic tests: although these do not reveal any observational artifacts, we cannot exclude the possibility that our detection is caused by the simulations not faithfully representing the statistical properties of the BOSS data. Our measurements can be used to constrain physical models of parity violation. As an example, we consider a coupling between the inflaton and a $U(1)$ gauge field and place bounds on the latter's energy density, which are several orders of magnitude stronger than those previously reported. Upcoming probes such as DESI and Euclid will reveal whether our detection of parity violation is due to new physics, and strengthen the bounds on a variety of models.

DOI: 10.1103/PhysRevD.106.063501

## I. INTRODUCTION

A detection of parity violation in cosmological observables would be a smoking gun for physics beyond the standard model, and it could provide crucial insights into the nature of dark matter, dark energy, and inflation. In the conventional paradigm, all cosmological correlators are symmetric under the parity operator $\mathbb{P}$, since gravity (along with all other standard model interactions except the weak force [1]) is $\mathbb{P}$-invariant. Despite this, a number of theoretical arguments suggest that parity-violating interactions *should* occur in the early universe, most notably to source baryogenesis. Creation of the current baryon asymmetry requires a process which violates charge and parity conservation [2,3]; a possible route is via leptogenesis, which, if sourced by gravity, must be parity violating, e.g., [4–8].

Additional sources of parity-violation include inflationary interactions between multiple fields, such as via the Chern-Simons term [9–14], generation of primordial magnetic fields [15–17], vector perturbations generated by cosmic strings or defects [18–20], reheating [21,22], Chern-Simons modified general relativity [11], and inflationary particle exchange [23,24], all of which leave distinctive imprints on late-time observables [9]. Potential evidence for such models was recently provided by [25], which found a $2.4\sigma$ hint (updated to $3.6\sigma$ in [26]) of parity violation in the cosmic microwave background (CMB). While some argue that this effect may be caused by interstellar dust emission [27] (though see [28,29]), it has nevertheless provided a resurgence of interest in these theories.

To constrain such phenomena, we require observables that are parity sensitive. Common choices are vector and tensor quantities, such as $B$ modes of the CMB [30], or those of galaxy ellipticities [31]. These satisfy $\mathbb{P}[B] = -B$, and can be combined in two-point correlators (e.g., $TB$ and $EB$ for the CMB, or $EB$ for weak lensing). Barring contamination by systematics, the observables should have no contribution from standard ΛCDM physics, but can be sourced by effects such as birefringence (whereupon the plane of the photon polarization is coherently rotated between the surface of last scattering and the observer, as in [25]), gravitational wave chirality [9,32–35], and multifield inflation [7,23,36]. Information is not limited to the two-point function, however; higher-order correlators

*ohep2@cantab.ac.uk

 

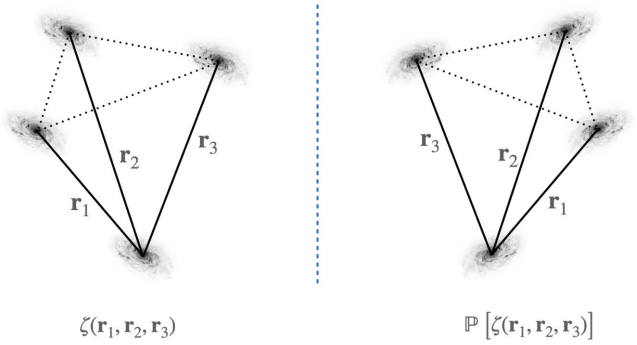


FIG. 1. Cartoon of the galaxy four-point correlation functions (4PCFs) considered in this work. In the left panel, we show the 4PCF, $\zeta(\boldsymbol{r}_1, \boldsymbol{r}_2, \boldsymbol{r}_3)$, which depends on the separation vectors of three secondary galaxies from a given primary. The right panel shows the parity-inverted 4PCF, $\mathbb{P}[\zeta(\boldsymbol{r}_1, \boldsymbol{r}_2, \boldsymbol{r}_3)]$, which corresponds to replacing $\boldsymbol{r}_i$ with $-\boldsymbol{r}_i$. Unlike for the 2PCF and 3PCF, the two configurations cannot be related by a rotation. The parity-even 4PCF is a sum of the two geometries (which have the same side lengths and relative angles), while the parity-odd 4PCF is a difference. In this work, the 4PCF is given as a function of three lengths ($r_1$, $r_2$, and $r_3$) and three internal angles (fixing the angle of the $\boldsymbol{r}_i$ vectors with the respect to the primary galaxy). The latter are represented by their harmonic-space momenta, $\ell_1$, $\ell_2$, and $\ell_3$, with odd-parity 4PCFs corresponding to odd $\ell_1 + \ell_2 + \ell_3$. Assuming standard ΛCDM physics, the two correlators shown in the figure should be equivalent, and thus the expectation value of the parity-odd 4PCF is zero.

such as *TTB* can give additional constraining power on effects such as birefringence [37].

When constructing observables from scalar fields (such as the galaxy density or CMB temperature), obtaining a parity-sensitive quantity is more difficult. As an example, the isotropic galaxy two-point correlation function (2PCF) is insensitive to parity, since the action of $\mathbb{P}$ is equivalent to a rotation, under which the statistic is invariant. In three dimensions, the isotropic $N$-point correlation functions (NPCFs) are parity sensitive only if $N > 3$; this applies also to the CMB, since the intrinsic fluctuations are the projection of a three-dimensional quantity. The simplest statistic with which to probe scalar parity violation is thus the four-point correlation function (4PCF), as pointed out in [23,38,39]. A cartoon of this is shown in Fig. 1.[1]

While a number of works have considered the 4PCF of the CMB [44,45] including its parity-odd contributions [38,46] (though only theoretically), the large scale structure (LSS) equivalent has been rarely explored. Given the influx of spectroscopic data expected in the next decade from DESI [47], Euclid [48], and Rubin [49], galaxy surveys seem to be a natural arena in which to hunt for parity-violating interactions, allowing constraints to exceed the CMB cosmic variance limit. Historically, use of the higher-point galaxy correlation functions has been hampered by the computational resources required for their estimation; naïvely, the 4PCF requires $\mathcal{O}(N_{\rm g}^4)$ operations to compute from $N_{\rm g}$ galaxies. Recent works have significantly improved upon this [50,51], with the algorithm of [51] requiring only $\mathcal{O}(N_{\rm g}^2)$ operations. This allows the 4PCF of current galaxy surveys to be computed in ∼30 CPU hours. The approach proceeds by first projecting the correlation function into a suitable angular basis [52]; thence, the integrals decouple and the 4PCF may be computed by summing over pairs of galaxies. This naturally generalizes to higher dimensions, as well as to anisotropic correlators [53]. As first pointed out in [52] there is a natural separation of the parity-even and parity-odd isotropic basis functions. The even-parity component can be used to place constraints on gravitational non-Gaussianity from a hitherto unexplored statistic [54]. The use of the parity-odd basis to measure parity violation in the galaxy four-point correlation was first proposed in [39] and is carried out in this work (see also [23]).

There are two main ways in which parity violation can be probed using the galaxy 4PCF. First, one may place constraints on the amplitudes of specific physical models given their associated theoretical predictions. This is an approach often used in the analysis of CMB three- and four-point functions, for example, in non-Gaussianity studies, which typically exploit separability of the underlying theoretical templates for significant computational gain [55]. This approach was also suggested in [23,46,56], and allows for targeted constraints on specific models of early universe particle exchange, via a search for their specific isotropy-violating signatures. An alternative method would be to first measure the *full* galaxy 4PCF in some set of bins, then perform a blind test, looking for the signatures of *any* physical model (and systematic effects). This approach is possible since the parity-odd 4PCF receives no contribution in ΛCDM, including from general relativistic and baryonic effects. Given the multitude of possible models for parity violation, we will principally adopt the second strategy in this work, though we demonstrate also the first, by placing constraints on a specific model involving Chern-Simons terms in the inflationary Lagrangian. Analysis using the galaxy 4PCF comes with its complexities, however. In particular, the high dimensionality of the statistic prohibits conventional mock-based $\chi^2$ analyses. To alleviate this, we include a data-compression step, facilitated using a theoretical model of the 4PCF covariance [57], which dramatically reduces the number of bins without introducing bias. It is not guaranteed that the 4PCF likelihood be Gaussian, however (see Appendix A and [58,59]); to provide a fully robust, yet conservative, test for parity violation, we make use of a likelihood-free inference technique, involving a suite of realistic simulations. We caution that such blind tests are

[1]Large scale structure correlators are sensitive also to redshift-space distortions [40,41], giving dependence of the statistic on line-of-sight velocities [42]. This enables vector-type parity violation to be probed in the 3PCF [43], though it requires careful modeling of galaxy velocities.

naturally subject to systematic uncertainties, some of which will be explored in this work. The results below represent the first constraints on scalar parity violation from LSS data.

The remainder of this work is structured as follows. In Sec. II, we present the parity-odd 4PCF estimator, including the corrections necessary to account for nonuniform survey geometry, before we discuss the data and covariance matrices in Sec. III. Analysis methods are considered in Sec. IV, with the corresponding constraints on parity violation presented in Sec. V. In Sec. VI, we include a number of systematic checks and a brief discussion of potential biases in the approach. Section VII discusses parity-breaking phenomena including the presentation of an inflationary model for the parity-odd 4PCF, based on a Chern-Simons coupling, whose amplitude is then bounded using the BOSS data. We conclude in Sec. VIII, with Appendixes A and B discussing the impacts of likelihood non-Gaussianity and sketching the derivation of the Chern-Simons 4PCF template. JUPYTER notebooks containing our analysis pipeline can be found on GitHub.[2]

*Note on Blinding.*—To limit confirmation bias, the BOSS data were sent to an external collaborator (M. König) after computation, and not revealed until the analysis pipeline was constructed and tested. The initial draft of the paper was also written before unblinding (encompassing all sections except Sec. VI and Appendix A), with the BOSS data replaced by that from a single mock dataset.

## II. MEASURING THE PARITY-ODD 4PCF

We begin by outlining our estimator for the parity-odd 4PCF, which is implemented in the public ENCORE code.[3] Further details of the formalism can be found in [51] (for the general NPCF estimator and ENCORE), in [52] (for the basis functions), in [54] (for the parity-even 4PCF), in [39] (for an overview of the parity-odd 4PCF), and in [53] (for extensions beyond flat three-dimensional [3D] space).

### A. Isotropic basis functions

Given a (scalar) density field $\delta(\boldsymbol{r})$, the 4PCF is defined as

$$\boxed{\zeta(\boldsymbol{r}_1,\boldsymbol{r}_2,\boldsymbol{r}_3) \equiv \langle \delta(\boldsymbol{s})\delta(\boldsymbol{s}+\boldsymbol{r}_1)\delta(\boldsymbol{s}+\boldsymbol{r}_2)\delta(\boldsymbol{s}+\boldsymbol{r}_3)\rangle,} \quad (1)$$

where $\langle\cdots\rangle$ represents a statistical average over realizations of $\delta$. A cartoon of this parametrization is shown in Fig. 1. By statistical homogeneity, the 4PCF is independent of the absolute coordinate $\boldsymbol{s}$.

As demonstrated in [51,53], a complete angular basis for the isotropic $N$-point correlation functions is given by the isotropic basis functions of $(N-1)$ coordinates defined in [52] (see also the TriPoSH formalism; [60]).[4] For $N=4$, the basis functions are

$$\mathcal{P}_{\ell_1\ell_2\ell_3}(\hat{\boldsymbol{r}}_1,\hat{\boldsymbol{r}}_2,\hat{\boldsymbol{r}}_3) \equiv (-1)^{\ell_1+\ell_2+\ell_3}\sum_{m_1m_2m_3}\begin{pmatrix}\ell_1 & \ell_2 & \ell_3\\ m_1 & m_2 & m_3\end{pmatrix} \times Y_{\ell_1m_1}(\hat{\boldsymbol{r}}_1)Y_{\ell_2m_2}(\hat{\boldsymbol{r}}_2)Y_{\ell_3m_3}(\hat{\boldsymbol{r}}_3), \quad (2)$$

where $Y_{\ell m}(\hat{\boldsymbol{r}})$ is a spherical harmonic, the $3\times 2$ matrix is a Wigner $3-j$ symbol, and the $m_i$ summations run over integer $m_i \in [-\ell_i,\ell_i]$. Such functions arise from the theory of angular momentum addition and are specified by three non-negative integers $\{\ell_1,\ell_2,\ell_3\}$, encoding the relative orientation of $\hat{\boldsymbol{r}}_1$, $\hat{\boldsymbol{r}}_2$, and $\hat{\boldsymbol{r}}_3$. Due to the $3-j$ symbol, the integers must obey the triangle condition $|\ell_1-\ell_2| \le \ell_3 \le \ell_1+\ell_2$, and we additionally enforce $\ell_i \le \ell_{\rm max}$. In practice, we restrict to relatively low $\ell_{\rm max}$, which gives an angular resolution of $\theta_{\rm min} \approx 2\pi/\ell_{\rm max}$ for the internal angles of the 4PCF tetrahedron. The basis functions have the following properties under parity and conjugation transformations (for parity operator $\mathbb{P}$):

$$\mathbb{P}[\mathcal{P}_{\ell_1\ell_2\ell_3}(\hat{\boldsymbol{r}}_1,\hat{\boldsymbol{r}}_2,\hat{\boldsymbol{r}}_3)] = (-1)^{\ell_1+\ell_2+\ell_3}\mathcal{P}_{\ell_1\ell_2\ell_3}(\hat{\boldsymbol{r}}_1,\hat{\boldsymbol{r}}_2,\hat{\boldsymbol{r}}_3),$$
$$\mathcal{P}^*_{\ell_1\ell_2\ell_3}(\hat{\boldsymbol{r}}_1,\hat{\boldsymbol{r}}_2,\hat{\boldsymbol{r}}_3) = (-1)^{\ell_1+\ell_2+\ell_3}\mathcal{P}_{\ell_1\ell_2\ell_3}(\hat{\boldsymbol{r}}_1,\hat{\boldsymbol{r}}_2,\hat{\boldsymbol{r}}_3), \quad (3)$$

implying that the basis is parity-odd and pure imaginary if $\ell_1+\ell_2+\ell_3$ is odd, and parity-even and real else. Furthermore, Eq. (2) is invariant under joint rotations of all three separation vectors, i.e., $\{\boldsymbol{r}_1,\boldsymbol{r}_2,\boldsymbol{r}_3\} \to \{R\boldsymbol{r}_1,R\boldsymbol{r}_2,R\boldsymbol{r}_3\}$, for arbitrary rotation matrix $R$.

The isotropic part of the galaxy 4PCF can be decomposed into the basis of (2),

$$\zeta_{\rm iso}(\boldsymbol{r}_1,\boldsymbol{r}_2,\boldsymbol{r}_3) = \sum_{\ell_1\ell_2\ell_3}\zeta_{\ell_1\ell_2\ell_3}(r_1,r_2,r_3)\mathcal{P}_{\ell_1\ell_2\ell_3}(\hat{\boldsymbol{r}}_1,\hat{\boldsymbol{r}}_2,\hat{\boldsymbol{r}}_3), \quad (4)$$

where the coefficients $\zeta_{\ell_1\ell_2\ell_3}$ (hereafter denoted "multiplets") can be obtained via the orthonormality of $\mathcal{P}_{\ell_1\ell_2\ell_3}$.[5] Given the transformation properties of (3), we find a natural split of $\zeta_{\rm iso}$ into parity-even and parity-odd parts:

$$\zeta_+(\boldsymbol{r}_1,\boldsymbol{r}_2,\boldsymbol{r}_3) = \sum_{\ell_1+\ell_2+\ell_3={\rm even}}\zeta_{\ell_1\ell_2\ell_3}(r_1,r_2,r_3)\mathcal{P}_{\ell_1\ell_2\ell_3}(\hat{\boldsymbol{r}}_1,\hat{\boldsymbol{r}}_2,\hat{\boldsymbol{r}}_3),$$
$$\zeta_-(\boldsymbol{r}_1,\boldsymbol{r}_2,\boldsymbol{r}_3) = \sum_{\ell_1+\ell_2+\ell_3={\rm odd}}\zeta_{\ell_1\ell_2\ell_3}(r_1,r_2,r_3)\mathcal{P}_{\ell_1\ell_2\ell_3}(\hat{\boldsymbol{r}}_1,\hat{\boldsymbol{r}}_2,\hat{\boldsymbol{r}}_3). \quad (5)$$

[2]Available at github.com/oliverphilcox/Parity-Odd-4PCF.
[3]Available at github.com/oliverphilcox/encore.
[4]The approach naturally extends to *anisotropic* correlators [53], though we do not consider them in this work.
[5]Since the anisotropic basis functions are orthogonal to those of (2), the decomposition in (4) holds regardless of whether the full statistic is isotropic.

These satisfy $\mathbb{P}[\zeta_\pm(\boldsymbol{r}_1,\boldsymbol{r}_2,\boldsymbol{r}_3)] = \pm\zeta_\pm(\boldsymbol{r}_1,\boldsymbol{r}_2,\boldsymbol{r}_3)$ and may be related to the sum and difference of the two panels in Fig. 1. In this work, we restrict to odd $\ell_1+\ell_2+\ell_3$, and thus consider the (purely imaginary) parity-odd 4PCF.

### B. 4PCF estimator

Invoking the ergodic principle, we may estimate the full 4PCF as an integral over four density fields,

$$\hat{\zeta}(\boldsymbol{r}_1,\boldsymbol{r}_2,\boldsymbol{r}_3) \equiv \frac{1}{V}\int d\boldsymbol{s}\,\delta(\boldsymbol{s})\delta(\boldsymbol{s}+\boldsymbol{r}_1)\delta(\boldsymbol{s}+\boldsymbol{r}_2)\delta(\boldsymbol{s}+\boldsymbol{r}_3), \quad (6)$$

where $V$ is the integration volume. This is unbiased; i.e., it has expectation $\mathbb{E}[\hat{\zeta}] = \zeta$. Since the basis functions of (2) are orthonormal [52], Eq. (6) can be used to construct an estimator for the 4PCF basis coefficients:

$$\begin{aligned}\hat{\zeta}_{\ell_1\ell_2\ell_3}(r_1,r_2,r_3) &= \int d\hat{\boldsymbol{r}}_1 d\hat{\boldsymbol{r}}_2 d\hat{\boldsymbol{r}}_3 \mathcal{P}^*_{\ell_1\ell_2\ell_3}(\hat{\boldsymbol{r}}_1,\hat{\boldsymbol{r}}_2,\hat{\boldsymbol{r}}_3)\hat{\zeta}(\boldsymbol{r}_1,\boldsymbol{r}_2,\boldsymbol{r}_3)\\ &= \frac{1}{V}\sum_{m_1m_2m_3}\begin{pmatrix}\ell_1 & \ell_2 & \ell_3\\ m_1 & m_2 & m_3\end{pmatrix}\\ &\quad\times\int d\boldsymbol{s}d\hat{\boldsymbol{r}}_1 d\hat{\boldsymbol{r}}_2 d\hat{\boldsymbol{r}}_3\delta(\boldsymbol{s})\delta(\boldsymbol{s}+\boldsymbol{r}_1)\delta(\boldsymbol{s}+\boldsymbol{r}_2)\delta(\boldsymbol{s}+\boldsymbol{r}_3)Y_{\ell_1m_1}(\hat{\boldsymbol{r}}_1)Y_{\ell_2m_2}(\hat{\boldsymbol{r}}_2)Y_{\ell_3m_3}(\hat{\boldsymbol{r}}_3),\end{aligned} \quad (7)$$

using the conjugate properties of (3). Defining the harmonic coefficients

$$\boxed{a_{\ell m}(\boldsymbol{s};r) \equiv \int d\boldsymbol{r}\delta(\boldsymbol{s}+\hat{\boldsymbol{r}})Y_{\ell m}(\hat{\boldsymbol{r}}),} \quad (8)$$

this is separable in $\hat{\boldsymbol{r}}_i$:

$$\boxed{\hat{\zeta}_{\ell_1\ell_2\ell_3}(r_1,r_2,r_3) = \sum_{m_1m_2m_3}\begin{pmatrix}\ell_1 & \ell_2 & \ell_3\\ m_1 & m_2 & m_3\end{pmatrix}\int\frac{d\boldsymbol{s}}{V}\delta(\boldsymbol{s})a_{\ell_1m_1}(\boldsymbol{s};r_1)a_{\ell_2m_2}(\boldsymbol{s};r_2)a_{\ell_3m_3}(\boldsymbol{s};r_3).} \quad (9)$$

For a discrete density field defined by $N_{\rm g}$ particles at positions $\{\boldsymbol{x}_i\}$ with weights $w_i$, the estimator can be written as a sum:

$$\begin{aligned}a_{\ell m}(\boldsymbol{x}_i;r) &\equiv \sum_{j=1}^{N_{\rm g}} w_j Y_{\ell m}(\widehat{\boldsymbol{x}_j-\boldsymbol{x}_i})\delta_{\rm D}(r-|\boldsymbol{x}_j-\boldsymbol{x}_i|),\\ \hat{\zeta}_{\ell_1\ell_2\ell_3}(r_1,r_2,r_3) &= \sum_{i=1}^{N_{\rm g}}\sum_{m_1m_2m_3}\begin{pmatrix}\ell_1 & \ell_2 & \ell_3\\ m_1 & m_2 & m_3\end{pmatrix}w_i a_{\ell_1m_1}(\boldsymbol{x}_i;r_1)a_{\ell_2m_2}(\boldsymbol{x}_i;r_2)a_{\ell_3m_3}(\boldsymbol{x}_i;r_3),\end{aligned} \quad (10)$$

where the Dirac delta, $\delta_{\rm D}$, ensures that we count only secondary particles $j$, separated from the primary $i$, by a distance $r$. Since we must compute $a_{\ell m}$ at the location of each primary particle, the estimator requires a sum over pairs of particles, and thus has complexity $\mathcal{O}(N_{\rm g}^2)$; in practice, the scaling is closer to linear, as the $m_i$ summation is rate limiting for large $\ell_{\rm max}$ [51]. By replacing the Dirac function in (10) by a binning function of finite width, the estimator extends to bin-averaged 4PCF estimates; we refer the reader to [51,54] for details. We further note that the 4PCF contains also a "disconnected" piece sourced by two copies of the 2PCF. While this can be subtracted at the estimator level directly [54], it does not contribute to parity-odd multiplets and will thus be ignored henceforth.

### C. Edge correction

Finally, estimator (9) must be modified to account for the nonuniform survey geometry. For this purpose, we first define the 4PCF using the generalized Landy-Szalay form [51,61,62]

$$\hat{\zeta}(\boldsymbol{r}_1,\boldsymbol{r}_2,\boldsymbol{r}_3) \equiv \frac{\mathcal{N}(\boldsymbol{r}_1,\boldsymbol{r}_2,\boldsymbol{r}_3)}{\mathcal{R}(\boldsymbol{r}_1,\boldsymbol{r}_2,\boldsymbol{r}_3)}, \quad (11)$$

where $\mathcal{N}$ and $\mathcal{R}$ are the 4PCF estimates obtained from "data-minus-random" and random catalogs, respectively, both of which are modulated by the survey window function. Following some algebra, the *edge-corrected* 4PCF multiplets are given by

$$\zeta_{\ell_1\ell_2\ell_3}(r_1,r_2,r_3) = \sum_{\ell_1'\ell_2'\ell_3'} [M^{-1}]^{\ell_1'\ell_2'\ell_3'}_{\ell_1\ell_2\ell_3}(r_1,r_2,r_3)\frac{\mathcal{N}_{\ell_1'\ell_2'\ell_3'}(r_1,r_2,r_3)}{\mathcal{R}_{000}(r_1,r_2,r_3)}, \tag{12}$$

defining the coupling matrix

$$M^{\ell_1'\ell_2'\ell_3'}_{\ell_1\ell_2\ell_3}(r_1,r_2,r_3) = \frac{(-1)^{\ell_1'+\ell_2'+\ell_3'}}{(4\pi)^{3/2}}\sum_{L_1L_2L_3}\frac{\mathcal{R}_{L_1L_2L_3}(r_1,r_2,r_3)}{\mathcal{R}_{000}(r_1,r_2,r_3)}\left[\prod_{i=1}^{3}\sqrt{(2\ell_i+1)(2L_i+1)(2\ell_i'+1)}\right]\begin{Bmatrix}\ell_1 & L_1 & \ell_1'\\ \ell_2 & L_2 & \ell_2'\\ \ell_3 & L_3 & \ell_3'\end{Bmatrix}$$
$$\times\begin{pmatrix}\ell_1 & L_1 & \ell_1'\\ 0 & 0 & 0\end{pmatrix}\begin{pmatrix}\ell_2 & L_2 & \ell_2'\\ 0 & 0 & 0\end{pmatrix}\begin{pmatrix}\ell_3 & L_3 & \ell_3'\\ 0 & 0 & 0\end{pmatrix}, \tag{13}$$

with the curly brackets indicating a Wigner $9-j$ symbol. This allows us to "undo" the effects of nonuniform survey geometry by measuring the 4PCF multiplets of the random field $\mathcal{R}$.[6] Note that there are two manners in which a parity-odd $\zeta$ can be sourced: parity-odd $\mathcal{N}$ and parity-even $\mathcal{R}$, or parity-odd $\mathcal{R}$ and parity-even $\mathcal{N}$.[7] For this reason, it is imperative to restrict to parity-odd multipets only *after* performing edge correction.

## III. DATA

### A. Data and simulations

Our dataset comprises galaxies from the twelfth data release (DR12) [64] of the Baryon Oscillation Spectroscopic Survey (BOSS), part of SDSS-III [65,66]. The survey contains two samples, CMASS and LOWZ, of which we use the former. This contains 587 071 (216 041) galaxies in the northern (southern) galactic cap (hereafter denoted NGC and SGC), across a redshift range $z \in [0.43, 0.7]$, and an effective redshift of $z_{\rm eff} = 0.57$.[8] We use a fiducial cosmology $\{\Omega_m = 0.31, \Omega_b h^2 = 0.022, h = 0.676, \sigma_8 = 0.8, n_s = 0.96, \sum m_\nu = 0.06\,\text{eV}\}$ to convert angles and redshifts to Cartesian coordinates (cf. [54,67]) and assign galaxy weights according to

$$w_{\rm tot} = (w_{\rm rf} + w_{\rm fc} - 1)w_{\rm sys}w_{\rm fkp}. \tag{14}$$

Here $w_{\rm rf}$, $w_{\rm fc}$, and $w_{\rm sys}$ correspond to redshift-failure, fiber-collision, and systematic weights, respectively, with $w_{\rm fkp} = [1 + n(z)P_0]^{-1}$ being the well-known Feldman-Kaiser-Peacock (FKP) weight [68] for background number density $n(z)$ and $P_0 = 10^4\ h^{-3}\,\text{Mpc}^3$. To model the survey geometry, we use the BOSS random catalogs, containing 50× more randoms than galaxies.

We additionally make use of a suite of $N_{\rm mocks} = 2048$ "MultiDark-Patchy" (hereafter PATCHY) simulations [69,70]. These are computed using an approximate gravity solver and calibrated to an $N$-body simulation, with halo occupation parameters adjusted such that the mocks well reproduce the BOSS two- and three-point statistics. These share the CMASS survey geometry and are assigned weights via

$$w_{\rm tot} = w_{\rm veto}w_{\rm fc}w_{\rm fkp}, \tag{15}$$

including the veto weight $w_{\rm veto}$. The mocks are generated with the parameter set $\{\Omega_m = 0.3071, \Omega_b h^2 = 0.02205, h = 0.6777, \sigma_8 = 0.8288, n_s = 0.96, \sum m_\nu = 0\ \text{eV}\}$ and coordinates are converted using the BOSS fiducial cosmology.

### B. 4PCF estimates

One of the main drawbacks with higher-order NPCFs is their dimensionality. To characterize the 4PCF, we must specify three multiplet indices $(\ell_1, \ell_2, \ell_3)$ and three radial bins $(r_1, r_2, r_3)$, which can lead to a statistic with a large number of (highly correlated) elements [51]. For this reason, we adopt a relatively coarse radial binning scheme using $N_r = 10$ linearly spaced radial bins in $[20, 160]\ h^{-1}\,\text{Mpc}$, giving $\Delta r = 14\ h^{-1}\,\text{Mpc}$. Furthermore, we enforce $r_2 > r_1 + \Delta r$ and $r_3 > r_2 + \Delta r$, to ensure that the separation between any two galaxies in the 4PCF tetrahedron is at least $\Delta r$ (cf. Fig. 1). This removes modes from the nonlinear region; these are difficult to model and can be strongly affected by baryonic physics. For the angular binning, we fix $\ell_{\rm max} = 5$, leading to a total of 56 radial bins and 111 multiplets (both parity-odd and parity-even), hence 6216 elements in the full 4PCF statistic. In the analysis of Sec. IV, we use only the 23 multiplets with odd $\ell_1 + \ell_2 + \ell_3$ and

[6]Note that this does not remove any geometry effects that couple to the *anisotropic* 4PCF, nor those coupling to the 4PCF multiplets with $\ell_i > L$, assuming an initial $\ell_{\rm max}$ of $L$. The former effect is expected to be small (and usually ignored for the 3PCF [63]), and the latter is ameliorated by discarding all multiplets containing $\ell_i = L$ after edge correction, justified by noting that the coupling matrix $M$ is close to tridiagonal.

[7]This occurs since the product of $3-j$ symbols in the coupling matrix is zero unless $\ell_1 + \ell_2 + \ell_3 + \ell_1' + \ell_2' + \ell_3' + L_1 + L_2 + L_3$ is even.

[8]Data are publicly available at data.sdss.org/sas/dr12/boss/lss/.

$\ell_i \leq 4$, giving a total of $N_\zeta = 1288$ elements; the rest are required for edge-correction (Sec. II C).

Computation of the 4PCF multiplets, $\zeta_{\ell_1\ell_2\ell_3}(r_1, r_2, r_3)$, is performed using the ENCORE code [51]. We separately measure the contributions from a random catalog and a set of 32 "data-minus-random" catalogs, each with 1.5× the galaxy density; the latter are averaged to form the $\mathcal{N}$ quantities entering the edge-correction equation (11), while the former give $\mathcal{R}$.[9] Using (12), the quantities are then combined to form the edge-corrected 4PCF multipoles.

For samples with similar number densities to BOSS, the runtime of ENCORE scales as $N_g N_r^3 (\ell_{\rm max}+1)^5$ [51], with computation dominated by the $m_i$ summations of (10) rather than estimation of the harmonic coefficients $a_{\ell m}$ (which scales as $N_g^2(1+\ell_{\rm max})^2$, albeit with a more modest prefactor). In practice, we parallelize computation using OPENMP, with each NGC (SGC) simulation requiring ∼32(6) CPU hours to analyze on a modern 16-core Intel processor, including edge correction. In total, analysis of the BOSS data and 2048 PATCHY mocks required ∼80 k CPU hours. This is comparable to the computational costs of the 2PCF analysis in Ref. [73], and is facilitated by the efficient nature of the ENCORE algorithm. We display a selection of the measured 4PCF multiplets in Fig. 2.

### C. Covariance matrices

The PATCHY mocks described in Sec. III A can be used to form a sample covariance of the 4PCF statistic in the standard manner:

$$\begin{aligned}\hat{\mathsf{C}}_{\ell_1\ell_2\ell_3;\ell_1'\ell_2'\ell_3'}(r_1, r_2, r_3; r_1', r_2', r_3') &= \frac{1}{N_{\rm mocks}-1}\sum_{i=1}^{N_{\rm mocks}} \left(\zeta^{(i)}_{\ell_1\ell_2\ell_3}(r_1, r_2, r_3) - \bar{\zeta}_{\ell_1\ell_2\ell_3}(r_1, r_2, r_3)\right) \\ &\times \left(\zeta^{(i)}_{\ell_1'\ell_2'\ell_3'}(r_1', r_2', r_3') - \bar{\zeta}_{\ell_1'\ell_2'\ell_3'}(r_1', r_2', r_3')\right), \end{aligned} \tag{16}$$

where $\zeta^{(i)}$ represents the $i$th 4PCF estimate (in the NGC or SGC region) and $\bar{\zeta}$ is the average over $N_{\rm mocks}$ realizations. Since the number of 4PCF bins exceeds the number of PATCHY mocks, this is not invertible, making it difficult to perform traditional $\chi^2$-based analyses. For this reason, we supplement the sample covariance with the analytic covariance described in [57]. Essentially, this computes:

$$\begin{aligned}{\rm Cov}(\boldsymbol{r}_1, \boldsymbol{r}_2, \boldsymbol{r}_3; \boldsymbol{r}_1', \boldsymbol{r}_2', \boldsymbol{r}_3') = &\int \frac{d\boldsymbol{s}}{V}\frac{d\boldsymbol{s}'}{V}\langle\delta(\boldsymbol{s})\delta(\boldsymbol{s}+\boldsymbol{r}_1)\delta(\boldsymbol{s}+\boldsymbol{r}_2)\delta(\boldsymbol{s}+\boldsymbol{r}_3)\delta(\boldsymbol{s}')\delta(\boldsymbol{s}'+\boldsymbol{r}_1')\delta(\boldsymbol{s}'+\boldsymbol{r}_2')\delta(\boldsymbol{s}'+\boldsymbol{r}_3')\rangle \\ &- \int \frac{d\boldsymbol{s}}{V}\langle\delta(\boldsymbol{s})\delta(\boldsymbol{s}+\boldsymbol{r}_1)\delta(\boldsymbol{s}+\boldsymbol{r}_2)\delta(\boldsymbol{s}+\boldsymbol{r}_3)\rangle\int\frac{d\boldsymbol{s}'}{V}\langle\delta(\boldsymbol{s}')\delta(\boldsymbol{s}'+\boldsymbol{r}_1')\delta(\boldsymbol{s}'+\boldsymbol{r}_2')\delta(\boldsymbol{s}'+\boldsymbol{r}_3')\rangle, \end{aligned} \tag{17}$$

where the statistical expectations can be expanded using Wick's theorem to yield products of four 2PCFs. The covariance is then projected into the angular basis of Sec. II A and simplified. The approach makes a number of assumptions:

(i) *Isotropy*: The 2PCF $\xi(\boldsymbol{r}) \equiv \langle\delta(\boldsymbol{s})\delta(\boldsymbol{s}+\boldsymbol{r})\rangle$ is assumed to be a function only of $|\boldsymbol{r}|$. This neglects redshift-space distortions, which have a nontrivial impact on the isotropic 4PCF covariance.

(ii) *Gaussianity*: The expectations entering (17) strictly contain additional contributions from higher-order correlators such as the 3PCF.

(iii) *Survey geometry*: While the 4PCF is edge corrected (Sec. II C), the same is not true for the covariance. The latter inherits nontrivial dependence on the survey geometry [74,75], which cannot be simply captured by modifying the survey volume or shot noise [57].

[9]If the algorithm's runtime scales as $N_g^2$, this partitioning minimizes the Poisson error at the fixed computational cost [71,72]. In our case, the scaling is closer to linear, and thus the total work is roughly independent of the partition size.

these reasons, we do *not* expect the analytic models of [57] to accurately predict the true covariance of BOSS. It is a relatively close approximation of matrix structure, however, and will thus be used as a proxy covariance to facilitate the analysis techniques described in Sec. IV. We construct the covariance using the same radial binning parameters as in Sec. III A, restricting to odd $\ell_1+\ell_2+\ell_3$. Following the prescription of [74] (but generalized to higher dimensions), we use an effective volume of 1.90 $h^{-3}\,{\rm Gpc}^3$ (0.77 $h^{-3}\,{\rm Gpc}^3$) and shot noise $P_{\rm shot} = 3130\ h^{-3}\,{\rm Mpc}^3$ (3160 $h^{-3}\,{\rm Mpc}^3$) for the NGC (SGC) subsample. The input 2PCFs are taken from a fit to the BOSS CMASS power spectrum, modeled using the effective field theory of large scale structure [76], as implemented in CLASS-PT [77].

Figure 3 compares the analytic and sample covariances for the NGC region, with the latter estimated from (16) using 2048 PATCHY mocks. Considering the correlation matrices [Fig. 3(a), defined as $\mathsf{R}_{ij} \equiv \mathsf{C}_{ij}/\sqrt{\mathsf{C}_{ii}\mathsf{C}_{jj}}$ for covariance $\mathsf{C}_{ij}$], we find good agreement between the two, indicating that the Gaussian theory model well reproduces the matrix structure. However, the diagonal elements [Fig. 3(b)] of the analytic covariance are roughly a factor of 2 less than those of the sample covariance. This is

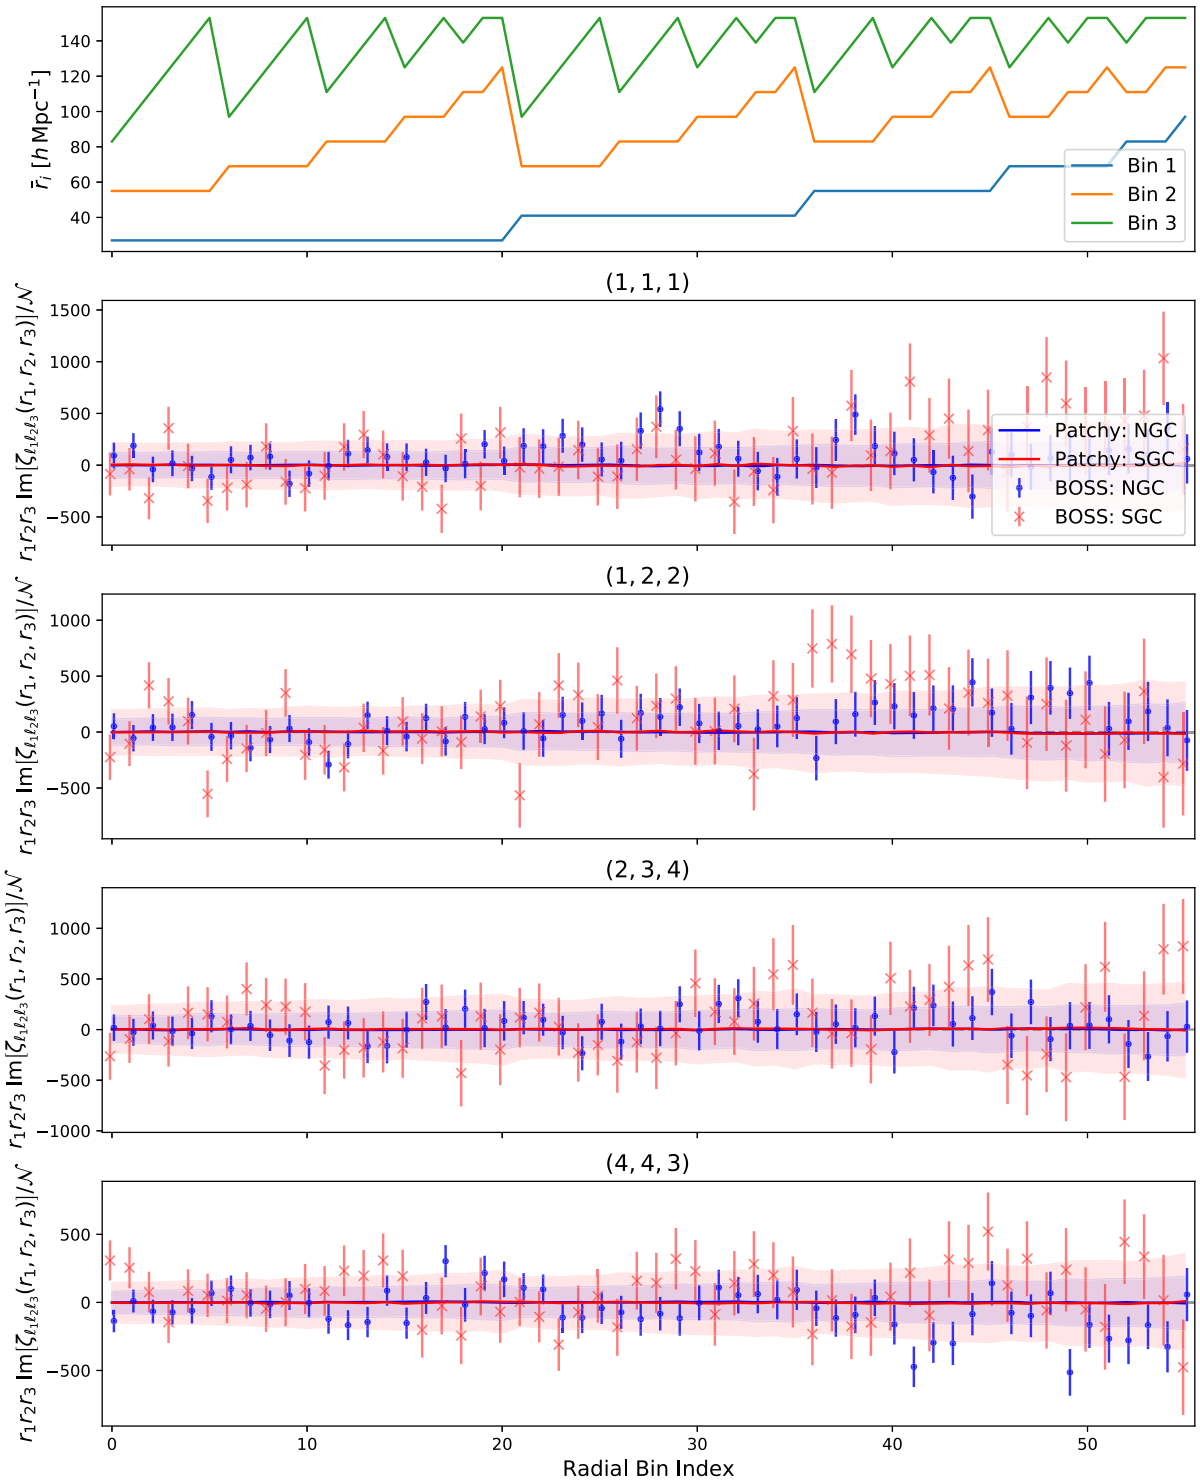


FIG. 2. Measurements of the parity-odd 4PCF from the BOSS CMASS galaxy sample, alongside those from 2048 PATCHY simulations. The NGC (SGC) results are shown in blue (red) bands, with the BOSS data shown as error bars, using the PATCHY variances. Results are displayed for a selection of $\{\ell_1, \ell_2, \ell_3\}$ multiplets (which specify the internal angles of the galaxy tetrahedron, as in Fig. 1), whose values are indicated by the title of each subfigure. In total, 23 parity-odd multiplets are included in the analysis of Sec. V. The horizontal axis specifies the radial bin combinations, $\{r_1, r_2, r_3\}$, with the central values of $r_1$, $r_2$, and $r_3$ in each bin shown in the top panel. These correspond to the distances of the secondary, tertiary, and quaternary galaxies from the primary in Fig. 1. For visibility, the 4PCF measurements are rescaled by a factor $-ir_1r_2r_3$. As expected, the PATCHY measurements show no signs of parity violation. Given the high correlation between neighboring bins, it is difficult to visually assess whether the BOSS dataset contains signatures of parity violation; this is quantified in Figs. 4 and 5.

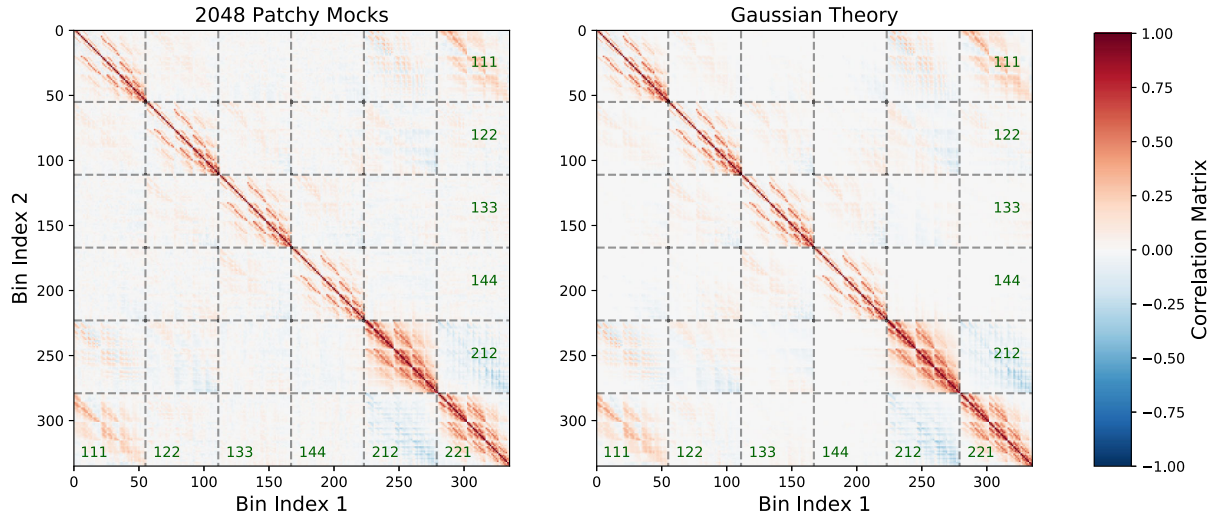


(a) Correlation Matrices

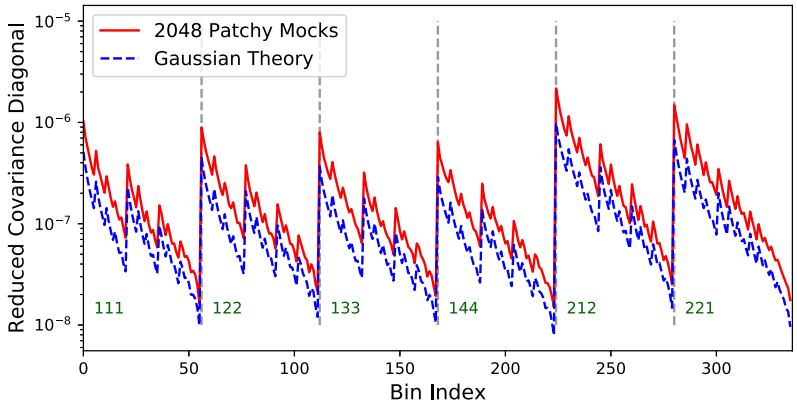


(b) Covariance Diagonal

FIG. 3. Comparison of the sample and analytic covariance matrices for the parity-odd 4PCF of the BOSS CMASS NGC region. The former are estimated using (16), while the latter use the approach of [57], which does not include redshift-space distortions, non-Gaussianity, or the effects of survey geometry. (a) Comparison of the correlation matrices (defined as the covariance matrices normalized by their diagonals); we see similar structure in both cases. The rows and columns represent the indices of the 4PCF, collapsed into one dimension, with each submatrix (indicated by the dotted lines) showing a different multiplet $\{\ell_1, \ell_2, \ell_3\}$, as labeled in green. Elements within a submatrix are ordered in increasing radii $r_1$, $r_2$, $r_3$. We include only the first six multiplets here; 23 are used in the analysis of Secs. V and VII C. (b) The corresponding diagonal elements of the covariance. Notably, the analytic covariance is an underestimate by a factor close to 2; we expect this to arise primarily due to the nonuniform survey geometry of the CMASS region [57].

likely to arise from the nontrivial survey geometry of the BOSS CMASS region [57] and prohibits direct use of the analytic covariance as a model for the 4PCF statistics.[10]

## IV. ANALYSIS METHODS

Below, we discuss two techniques that will be used to search for a signature of parity violation in Sec. V: (1) a nonparametric rank test, which does not require the likelihood to be Gaussian, and (2) data compression followed by a mock-based $\chi^2$ analysis. Both approaches make use of the smooth (but inaccurate) covariance matrix model of Sec. III C to overcome the difficulties associated with the high dimensionality of the 4PCF. To avoid confirmation bias, the pipeline implementing these techniques[11] was constructed before the BOSS data were unblinded.

### A. Nonparametric rank test

Nonparametric tests provide a powerful way to analyze data when the underlying likelihood is not known. Here, we

[10]Note that this discrepancy is *not* fully resolved by rescaling the theory covariance by a constant factor.

[11]Available at github.com/oliverphilcox/Parity-Odd-4PCF.

consider a *rank test*, examining the null hypothesis of zero parity-odd 4PCF. To implement this, we first define a test statistic, computed on both the data and a set of mocks. These simulations are required to obey the null hypothesis (i.e., be parity invariant) and have realistic noise properties. The test statistic measured from data is then compared to the empirical distribution obtained from the mocks, allowing construction of a detection significance. For example, if the data statistic exceeds that of 95% of the mocks, we may reject the null hypothesis at 95% C.L. The principal advantage of this approach is that it does not require a theoretical PDF for the test statistic; i.e., we do not have to assume the 4PCF to be a draw from some multivariate Gaussian. Indeed, the observed 4PCF does *not* appear to be Gaussian; this is explored in Appendix A. A limitation of such rank tests is that one cannot claim a detection at high significance; rather the maximal confidence level is $(1 - 1/N_{\rm mocks})$.

Below, we will use the following test statistic, dubbed the *pseudo-$\chi^2$*:

$$\tilde{\chi}^2 \equiv [\zeta^T \tilde{\mathsf{C}}^{-1} \zeta]_{\rm NGC} + [\zeta^T \tilde{\mathsf{C}}^{-1} \zeta]_{\rm SGC}, \tag{18}$$

where $\zeta$ is the set of measured parity-odd 4PCF multipoles (treated as a $N_\zeta$-dimensional vector) and $\tilde{\mathsf{C}}$ is the theoretical covariance matrix (Sec. III C). If $\tilde{\mathsf{C}}$ is equal to the sample covariance (in the limit of infinite mocks), Eq. (18) reduces to the usual $\chi^2$ statistic, given a fiducial model of zero parity-odd 4PCF and assuming the NGC and SGC regions to be independent. While the covariances are not quite equal in practice (Fig. 3), we expect (18) to produce a close-to-optimal weighting for the data, particularly if the likelihood is close to Gaussian. Furthermore, since the pseudo-$\chi^2$ statistic does not subtract off a mean, the rank test will naturally account for any spurious parity-odd contributions that are present in both PATCHY and BOSS. These might arise from imperfections in the edge-correction routine or light cone projection effects. To perform the test, we simply compute $\tilde{\chi}^2$ for BOSS and each of the $N_{\rm mocks} = 2048$ PATCHY simulations (Sec. III A), before assigning a detection significance from the empirical PATCHY PDF.

### B. Compressed Gaussian analysis

A common trick when dealing with high-dimensional statistics is to apply some form of data compression [e.g., [78–81]]. In general, this proceeds by projecting the data onto some (small) set of basis vectors, thus greatly reducing the dimensionality. When performing parameter inference, basis vectors are usually chosen to preserve the Fisher information matrix [78,80] or the log-likelihood [81]. Since our primary goal in this work is to search for signatures of parity violation in a model-agnostic fashion, we adopt a somewhat different compression scheme, following [54,79].

Here, we project the 4PCF onto a basis given by the eigenvectors of the theoretical covariance matrix (Sec. III C). Explicitly, we define the projected statistic

$$v \equiv \mathsf{U}^T \zeta, \tag{19}$$

where the orthogonal matrix $\mathsf{U}$ is specified by $\tilde{\mathsf{C}} = \mathsf{U}\Lambda\mathsf{U}^T$ for diagonal eigenvalue matrix $\Lambda$. The compressed statistic has covariance $\mathbb{E}[vv^T] = \mathsf{U}^T\mathsf{C}\mathsf{U}$, where $\mathsf{C}$ is the covariance of $\zeta$; if the theory and analytic covariances agree, this is diagonal and equal to $\Lambda$. In practice, we expect the compressed coefficients to be almost independent.

To perform dimensionality reduction, we must restrict to a subset of the aforementioned basis vectors. Given that we have no prior on the shape of a parity-violating 4PCF signal, we cannot select the basis vectors based on signal-to-noise considerations (as in [54,79]). Instead, we use the $N_{\rm eig}$ eigenvectors with smallest $\Lambda_i$, corresponding to the directions that can be most well measured.[12] This highlights the benefits of using the *theoretical* covariance matrix to perform the projection; since the sample covariance does not have full rank, its smallest eigenvalues are not well defined.

Following the selection of the basis vectors, we project both the BOSS data and the PATCHY mocks into the $N_{\rm eig}$-dimensional subspace using (19). As in (16), we can form a sample covariance for $v$ from the PATCHY measurements:

$$\hat{\mathsf{C}}_{v,\alpha\beta} = \frac{1}{N_{\rm mocks} - 1} \sum_{i=1}^{N_{\rm mocks}} (v_\alpha^{(i)} - \bar{v}_\alpha)(v_\beta^{(i)} - \bar{v}_\beta), \tag{20}$$

where $v_\alpha^{(i)}$ indicates the compressed 4PCF of the $i$th mock, with $\alpha, \beta \in \{1, \ldots, N_{\rm eig}\}$. Assuming $N_{\rm mocks} > N_{\rm eig}$, the sample covariance has full rank (unlike the uncompressed 4PCF covariance) and is thus invertible. In the low-dimensional subspace, analysis centers around the following statistic:

$$\hat{T}^2 = v^T \hat{\mathsf{C}}_v^{-1} v, \tag{21}$$

where we have assumed zero mean, as in the null hypothesis. If $v$ is assumed to be Gaussian distributed (a fair assumption if the dimensionality is small), $\hat{T}^2$ follows a $\chi^2$ distribution with $N_{\rm eig}$ degrees of freedom in the limit of large $N_{\rm mocks}$. In practice, we must account for noise in the sample covariance $\hat{\mathsf{C}}_v$. An approach is to add the "Hartlap" correction factor [82,83], leading to the modified statistic

$$\hat{H}^2 = f_H \times v^T \hat{\mathsf{C}}_v^{-1} v, \qquad f_H = \frac{N_{\rm mocks} - N_{\rm eig} - 2}{N_{\rm mocks} - 1}, \tag{22}$$

[12] Additional choices of basis functions can be found in Sec. VI D.

whose expectation is $\chi^2$. $\hat{H}^2$ is then analyzed using a $\chi^2$ distribution, assuming Gaussianity. However, this does not correctly treat the sample covariance noise and results in a PDF which is too sharply peaked if $N_{\rm eig}$ is close to $N_{\rm mocks}$ [84]. Instead, one should analyze the $\hat{T}^2$ statistic (21) directly, using the PDF:

$$f_T(T^2;n,p) = \frac{\Gamma[(n+1)/2]}{\Gamma(p/2)\Gamma[(n-p+1)/2]} \frac{n^{-p/2}(T^2)^{p/2-1}}{(T^2/n+1)^{(n+1)/2}}, \tag{23}$$

where $n = N_{\rm mocks} - 1$, $p = N_{\rm eig}$, and $\Gamma$ is the Gamma function [84]. When dealing with multiple independent datasets (i.e., the NGC and SGC 4PCF measurements), one may sum the two $T^2$ estimates; the resulting PDF is the convolution of two copies of (23) and is easily evaluated with a fast Fourier transform. This approach will be adopted for the main analysis of Sec. V to ensure that we do not claim a false detection of parity violation.

Finally, we comment on the validity of our compression scheme. By the Eckart-Young theorem [85], the scheme is optimal (in terms of inverse variance) in the limit of $\tilde{\mathsf{C}} = \mathsf{C}$ and a Gaussian covariance. Since the data and mocks are compressed in the same manner, it is unbiased for any choice of projection matrix $\mathsf{U}$ or dimension $N_{\rm eig}$. If too few basis vectors are used or if the theory covariance is far from the truth, the penalty (in the limit of large $N_{\rm mocks}$) is simply a reduced detection significance.[13] When using finite $N_{\rm mocks}$, increasing $N_{\rm eig}$ will also lead to increased noise in the covariance matrix $\hat{\mathsf{C}}_\nu$, somewhat lessening the detection significance. To incorporate these effects, a range of $N_{\rm eig}$ values will be considered in Sec. V.

## V. NULL-TEST RESULTS

We now proceed to assess whether the BOSS data contain signatures of parity violation. First, we consider the raw 4PCF measurements, displayed in Fig. 2 for a selection of multiplets $\{\ell_1, \ell_2, \ell_3\}$. As expected, the mean parity-odd 4PCF of PATCHY mocks appears close to zero. This functions as a useful consistency test for the analysis pipeline; errors in the 4PCF computation could have led to a detectable signal in PATCHY mocks, especially given that the parity-even 4PCF is large [54]. For the standard deviations, we find a rough scaling of $(r_1 r_2 r_3)^{-1}$, with enhanced amplitudes found for the SGC region due to its smaller volume (Sec. III). Moving to the BOSS results, we find considerable (highly correlated) scatter around zero, but no obvious signatures of parity violation.

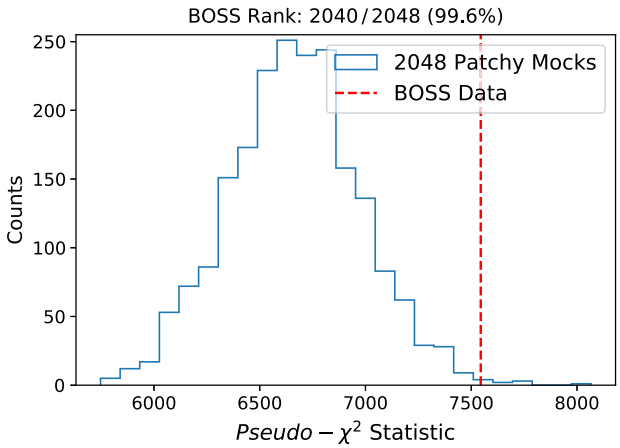


FIG. 4. Empirical distribution of the *pseudo*-$\chi^2$ test statistic defined in (18) from 2048 PATCHY mocks and the BOSS data. This is a proxy for the standard $\chi^2$ parameter, but uses the theoretical covariance matrix of Sec. III C. The data (shown as a vertical dashed line) have a Cumulative distribution function (CDF) of 99.6%; this is inconsistent with the null hypothesis of parity invariance at $2.9\sigma$. Note that this test does not assume Gaussianity of the likelihood, and naturally encompasses any spurious parity-odd contributions appearing in both the PATCHY mocks and the BOSS data. The significant detection of parity violation found by this test indicates either parity-violating physics or unresolved systematics. The above plot represents the main result of this work.

To examine this further, we turn to the nonparametric rank test of Sec. IV A. In Fig. 4, we plot the empirical distribution of the *pseudo*-$\chi^2$ test statistic; this appears to have broader tails than a standard $\chi^2$ distribution (most likely due to imperfections in the theoretical covariance $\tilde{\mathsf{C}}$), highlighting the importance of a nonparametric treatment. The BOSS data have a rank of 2040/2048, and an associated detection probability of 99.6% (equivalent to $2.9\sigma$). This is inconsistent with a random draw from the empirical distribution of PATCHY mocks and gives evidence for parity violation. A greater number of simulations would be needed to probe this detection at higher significance.

An additional test is given by the projected $\chi^2$-based analysis discussed in Sec. IV B. While this assumes a Gaussian distribution for the compressed statistic (somewhat justified by the reduced number of bins), it uses the sample covariance rather than the theoretical covariance, and thus optimally weights the data. In the top panel of Fig. 5, we show the theoretical and empirical distributions of the projected sample statistics $T^2$ and $H^2$ (Sec. IV B) from PATCHY, with the latter obtained via bootstrapping. For small $N_{\rm eig}$, the empirical distributions of both statistics

[13] This is easiest to show by considering the average $\chi^2$ difference between some signal $\zeta_0$ and the null hypothesis of $\mathbb{E}[\hat{\zeta}] = 0$. Without compression, $\Delta\chi^2 = \zeta_0^T \mathsf{C}^{-1} \zeta_0$, while following projection by some $N_\zeta \times N_{\rm eig}$ matrix $\mathsf{U}$, $\Delta\chi^2_{\rm proj} = \zeta_0^T \mathsf{U}(\mathsf{U}^T\mathsf{C}\mathsf{U})^{-1}\mathsf{U}^T\zeta_0$. If the projection is optimal, i.e., if $\mathsf{U}$ is the eigenvector matrix of $\mathsf{C}$, then $\Delta\chi^2_{\rm proj,opt} = \sum_{i=1}^{N_{\rm eig}} \bar{\zeta}^2_{0,i}/\Lambda_i$ where $\bar{\zeta}_0 = \mathsf{U}^T\zeta_0$. Since $\bar{\zeta}^2_{0,i} \geq 0$ and $\Lambda_i > 0$, it is clear that $\Delta\chi^2_{\rm proj,opt} \leq \Delta\chi^2$, with equality iff $N_{\rm eig} = N_\zeta$. A similar result holds in the more general case due to the properties of projection matrices; this is easily shown by rotating to a diagonal basis.

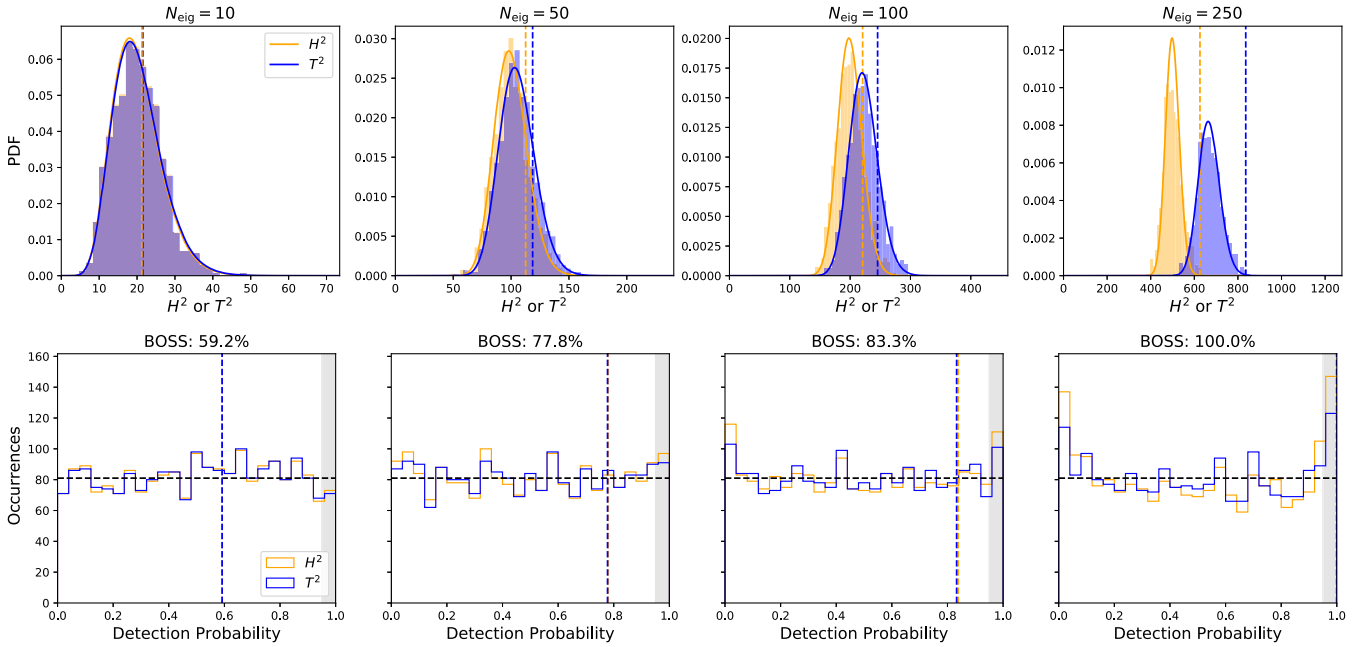


FIG. 5. Distributions of the $H^2$ and $T^2$ statistics (22) and (21) for the compressed BOSS data and PATCHY simulations. We show results using various numbers of basis vectors, $N_{\rm eig}$, as indicated by the titles, and note that the statistic includes both NGC and SGC measurements. The top panels compare the theoretical and empirical PDFs for each statistic, while the bottom panels display the CDFs. Results for BOSS are shown as vertical dashed lines in both cases. To compute the empirical distributions (shown as histograms), we apply bootstrapping; $(N_{\rm mocks} - 1)$ mocks are used to define a sample covariance, allowing computation of $H^2$ and $T^2$ for the excluded mock. The theoretical PDFs for $H^2$ and $T^2$ are the convolution of two $\chi^2$ or $T^2$ distributions (23). For BOSS, we report detection probabilities of 59.2%, 77.8%, 83.3%, and 100.0% from the $T^2$ statistic using $N_{\rm eig} = 10$, 50, 100, and 250 basis vectors, respectively. As in Fig. 4, we find mild evidence for parity violation, particularly as $N_{\rm eig}$ increases. We caution that this test assumes a Gaussian likelihood, which may lead to overestimated detection probabilities at high $N_{\rm eig}$ (as suggested by the somewhat skewed empirical distribution of PATCHY mocks at $N_{\rm eig} = 250$, and the results of Appendix A, which correspond to $N_{\rm eig} = N_\zeta$).

seem well-fit by their theory models, which is expected since $N_{\rm mocks}$ is considerably larger than $N_{\rm eig}$. At larger $N_{\rm eig}$ we begin to see discrepancies between the $H^2$ statistic and the accompanying $\chi^2$ theory model, with the former having a slightly narrower distribution. Considering the detection CDFs shown in the lower panel of Fig. 5, the distribution of PATCHY mocks appear somewhat nonuniform for the $H^2$ statistic at $N_{\rm eig} > 100$. In particular, 152 (199) mocks lie in the outermost 5% of the theory distribution for $N_{\rm eig} = 100$ ($N_{\rm eig} = 250$), in contrast to the $102 \pm 10$ expected. This echoes the conclusion of [84]; if $N_{\rm eig}$ is close to $N_{\rm mocks}$, improper treatment of covariance matrix noise may be dangerous and could lead to false detections of parity violation. For the $T^2$ distribution (which correctly treats such effects) we find somewhat better agreement, with 135 and 162 mocks in the outer 5% region, respectively. However, the distribution still appears to be somewhat skewed. We attribute this to intrinsic non-Gaussianity of the 4PCF likelihood (see [59] and Appendix A), whose effect increases as the size of the data vector increases, and the central limit theorem becomes less applicable.

For the BOSS data, the compressed Gaussian analysis gives detection significances of 59.2%, 77.8%, 83.3%, and 100.0% for $N_{\rm eig} = 10$, 50, 100, and 250, respectively, equivalent to $1.3\sigma$, $1.7\sigma$, $1.9\sigma$, and $3.9\sigma$ in a two-tail test. The fact that these results are a strong function of $N_{\rm eig}$ suggests that our projection scheme is inefficient, i.e., that 4PCF components dropped from the analysis carry significant information regarding the parity-violating signature.[14] This is a consequence of performing a blind analysis; given the lack of a physical model, we cannot choose basis vectors which maximize the signal to noise (though see Sec. VI D for results using alternative choices of the compression scheme). Overall, we find a weak preference for a non-Gaussian signal using the compressed analysis. This notwithstanding, we again find a preference for a nonzero parity-odd signal from the sample, though again caution that the $N_{\rm eig} = 250$ result may be artificially enhanced by likelihood non-Gaussianity.

[14]This differs from the conclusion of [54], which used a similar compression scheme to analyze the parity-even 4PCF. In the former work, basis vectors were chosen based on the mean non-Gaussian signal from the mocks, ensuring optimal linear compression. This is not possible in our case, since the mocks conserve parity.

To close, we comment on the implications of the results found herein. First, we note that the results are largely consistent between the two tests, with both finding a detection of large-scale parity violation at $\sim 3\sigma$. This cannot be caused by an incorrectly modeled likelihood (evidenced by the rank test), nor is it a result of our analysis incorrectly treating the window function (which would have led to the compressed analysis showing a parity-violating signature in the PATCHY mocks). This leaves two explanations: (1) new physics, and (2) unexplained systematics. While the first is a distinct possibility (and is not ruled out by other observations, since no former experiment has measured the 4PCF in a model-agnostic fashion), the second should also be carefully considered. A brief discussion of this is presented in the next section.

## VI. SYSTEMATIC TESTS

Below, we report the results of various checks performed to test the results of Sec. V, utilizing data cuts, mock catalogs, and rescaled statistics. In the final section, we will also comment on potential sources of systematic effects that could lead to a false detection of parity violation. All the analyses below were devised post-unblinding.

### A. Data partitioning

A simple test for systematic errors is to repeat the nonparametric rank test described in Sec. IV A for the two BOSS regions (NGC and SGC) separately. The two patches are in opposite hemispheres, have different calibrations, and are of different angular sizes; thus, this is a sensitive test of the effects of survey calibration and large-scale modes.

Figure 6 displays the *pseudo*-$\chi^2$ statistic for the two BOSS regions separately (plotting both the observational data and the empirical distribution from 2048 PATCHY mocks). Here, we find ranks of 2023/2048 and 1963/2048 for the NGC and SGC regions separately, with a lower significance found for the smaller-volume region,

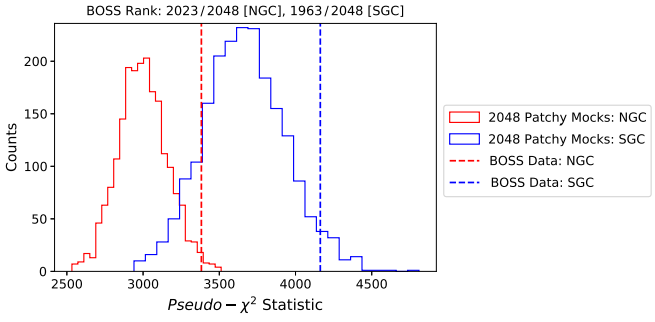


FIG. 6. As Fig. 4, but analyzing the two observational regions (NGC and SGC) separately. The NGC region (red) contains almost 3 times as many galaxies and is thus expected to be more constraining. We report detection significances of 98.8% (95.8%) for NGC (SGC), corresponding to $2.5\sigma$ ($2.0\sigma$).

as expected. Neither value is enough to claim a detection on its own; however, the trends from the two are in agreement, and their combination reaches the $2.9\sigma$ level reported in Sec. III. This test does not reveal any clear differences between the two regions; the results are consistent with our expectations if this is indeed a *bona fide* detection of parity violation.

### B. Dependence on radial and angular scales

An additional test is to examine the impact of radial and angular binning on the detection significance reported in Sec. V. For this purpose we will primarily use the rank test of Sec. IV A, and adjust the criteria for which bins are included in the analysis, always including a subset of those discussed in Sec. III B. In general, we expect the detection significance to decrease somewhat as the dimensionality (and hence number of measured Fourier modes) is reduced, relative to the initial size of $N_\zeta = 1288$). This will be particularly apparent if we excise regions in which the signal is strongest.

First, we consider changing the radial binning strategy. Two variations are possible: we can filter based on the distances of secondary galaxies from the primary or those of the secondary galaxies from each other (cf. Fig. 1). To test the former, we will separately remove the first and last radial bins in each dimension, i.e., that with $r_i \in [20, 34]\ h^{-1}\,\mathrm{Mpc}$ and $[146, 160]\ h^{-1}\,\mathrm{Mpc}$, respectively. When the minimum radius is increased, we have a modest decrease in dimensionality (to $N_\zeta = 805$) and a detection rank of 1949/2048 ($2.0\sigma$). If the maximum radius is instead reduced, $N_\zeta$ decreases by a third, but we do not find a change in the overall detection significance. The first result is unsurprising: the variance scales approximately as $(r_1 r_2 r_3)^{-2}$; thus, if the signal has support over a range of radii, removing the low-$r$ bins would reduce the detection significance. In the second case, the lack of variation suggests that the detection is not caused by some spurious ultralarge mode (arising from foregrounds, for example), though we caution that the signal to noise in these modes is the smallest.

To vary the allowed distances between secondary galaxies, we may modify the restrictions on the allowed tetrahedral shapes. In particular, accepting configurations with $r_1 < r_2 + \gamma\Delta r < r_3 + 2\gamma\Delta r$ restricts the internal separations to be at least $\gamma\Delta r$ (with $\gamma = 1$ in the fiducial analysis of Sec. V). If we set $\gamma = 0$, and thus include all bins satisfying $r_1 < r_2 < r_3$, secondary galaxies can be arbitrarily close together, and the size of the data vector increases to $N_\zeta = 2760$. If the signal contains small-scale power, this should increase the detection significance; here, the nonparametric test gives a rank of 2045/2048 (almost at the saturation point), equivalent to $3.2\sigma$, with the $N_{\mathrm{eig}} = 100$ (250) compressed $\chi^2$-based analysis giving a detection probability of 99.7% (100.0%), equivalent to $3.4\sigma$ ($4.3\sigma$).

This suggests that small-scale modes are of importance, though we note that they could contain contributions from parity-breaking physics on halo scales, such as magneto-hydrodynamics (whose contributions are generally expected to be small in the main analysis, since we restrict to separations above 14 $h^{-1}$ Mpc). If we instead fix $\gamma = 2$, forcing secondary galaxies to be separated by at least 28 $h^{-1}$ Mpc, we find a sharp reduction in dimensionality to $N_\zeta = 460$, but only a slight decrease in the detection rank (to 1970/2048, or $2.1\sigma$). This suggests that the signal causing the detection contains nontrivial support on relatively large (and generally well understood) scales.

Next, we consider varying the angular binning, by altering the maximum multipole $\ell_{\rm max}$, and thus the values of $\{\ell_1, \ell_2, \ell_3\}$ (Sec. II) that are used in the analysis. Setting $\ell_{\rm max} = 3$ reduces the dimensionality to $N_\zeta = 616$ and changes the detection significance to 1970/2048, or 96.2% ($2.1\sigma$). The physical action of this is to reduce the number of squeezed tetrahedra in the analysis (which have smaller internal tetrahedron angles, i.e., large $\ell$). If $\ell_{\rm max}$ is further reduced to 2 (with $N_\zeta = 224$) the rank falls to 1466/2048, or 71.6% ($1.1\sigma$). Our conclusions for $\ell_{\rm max} = 3$ are similar to the above: the signal contains contributions from small-scale modes, but is not entirely dominated by them, while the $\ell_{\rm max} = 2$ result will occur primarily due to the much reduced dimensionality, and thus signal to noise. We further note that the multiplets with the greatest impact on the detection significance are $\{\ell_1, \ell_2, \ell_3\} = \{1, 1, 1\}$, $\{1, 2, 2\}$ and permutations thereof. This is again unsurprising, given that these are the modes with the greatest signal to noise, but may serve to indicate that any signal observed is not entirely from some squeezed limit.

### C. Realization-dependent rescaling

Under null assumptions (and assuming Gaussianity for the sake of illustration), the rank test compares the variance of the observational data with that of the PATCHY mocks. As such, a spurious detection of parity violation could be caused by the simulations underestimating the true covariance, for example, due to mismatches in the galaxy bias parameters or the underlying cosmology. From Fig. 4, we see that excellent agreement between data and mocks can be obtained if one inflates the covariance of the PATCHY simulations by 13%. Under the assumption of a Gaussian 4PCF covariance, this requires the 2PCF to be rescaled by 3%. While this may seem straightforward, it is somewhat nontrivial, since the PATCHY simulations are calibrated to match the BOSS two- and three-point clustering statistics on small scales [69,70].

To probe this, we consider normalizing the odd-parity 4PCF measurements by their even-parity (more specifically, disconnected) counterparts and reapplying the rank test of Sec. IV A. To perform this robustly, we first fit the Gaussian 4PCF contribution to a simple theory template (presented in [86]), extracting a single overall amplitude for each realization.[15] The odd-parity 4PCF measurements (for both the BOSS data and each PATCHY mock) are then divided by this amplitude, which will remove the dominant effect of an overall rescaling of the PATCHY covariance compared to that of BOSS. While the PATCHY covariance need not be wrong by a simple constant, this is expected to capture the leading-order effect, and particularly accounts for the unknown value of $b_1\sigma_8$. We note that this prescription cannot be easily applied to the compressed Gaussian analyses, since it violates the Gaussian assumption on large scales due to sample-variance cancellation.

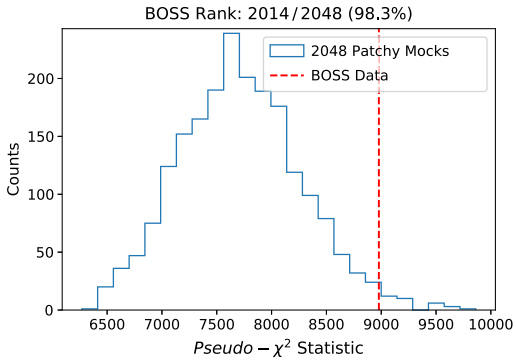


FIG. 7. As Fig. 4, but normalizing the 4PCF measurements by a realization-dependent amplitude, as described in Sec. VI C. This will account for an overall rescaling factor between the BOSS data and PATCHY simulations, arising, for example, from a different value of $\sigma_8$. In this case, the detection significance is slightly reduced; however, we still find hints of parity violation at the $2.4\sigma$ level.

In practice, we find an NGC (SGC) normalization factor of 0.85 (0.99) for BOSS, and $0.92 \pm 0.04$ ($0.94 \pm 0.06$) for the 2048 PATCHY simulations. Notably, the BOSS NGC data have a smaller factor; this will shift the corresponding *pseudo*-$\chi^2$ value of BOSS toward those of the mocks. In Fig. 7, we show the corresponding rank test results, finding that the detection significance is slightly reduced (to 2014/2048), matching our expectation. However, this is consistent with the broadened posterior associated with this test (comparing the widths of the empirical distributions in Figs. 4 and 7), and we still find a weak detection of parity violation, now at $2.4\sigma$. It is clear that the detection cannot be entirely removed by a simple rescaling; if differences in the statistical properties of simulations and observational data are to blame, they must be scale dependent.

### D. Choice of compression scheme

For the compressed Gaussian analysis of Sec. IV B, it is important to choose a set of basis vectors that allow for

[15] An alternative approach is to divide the odd-parity 4PCF by the value of the disconnected (i.e., Gaussian) 4PCF in each bin; this is not performed since the latter statistic suffers from significant cosmic variance on large scales and can be negative.

significant dimensionality reduction while preserving key features of the data. For null tests such as that of Sec. IV B, this is nontrivial, since we do not have prior knowledge of the form of a parity-violating model. For this reason, the analysis of Sec. III adopted a set of basis vectors selected using a minimum variance criterion; this is equivalent to maximizing signal to noise assuming a uniform signal in all bins.

An alternative approach would be to assert some typical form for the parity-violating signal, and use this to select a sensible set of basis vectors onto which the measured 4PCFs are projected. A simple choice is $\zeta_\Lambda(r_1, r_2, r_3) \propto (r_1 r_2 r_3)^{-1}$ (matching the approximate scaling of the error bars). In this case, the compressed Gaussian analysis gives detection significances of $0.8\sigma$, $1.8\sigma$, $3.5\sigma$ for $N_{\rm eig} = 50$, 100, 250, respectively. A more nuanced choice would be to weight by the inflationary parity-breaking 4PCF model introduced in Sec. VII. This has a physically reasonable form (albeit specific to a single parity-breaking phenomenon) and leads to Gaussian significances of $0.7\sigma$, $1.5\sigma$, $2.3\sigma$, respectively. Finally, we consider fixing the fiducial model to the $\{\ell_1, \ell_2, \ell_3\} = \{0, 0, 0\}$ multiplet of the disconnected 4PCF. This will indicate whether the observed signal was sourced by incomplete subtraction of the disconnected component. In this case, we find detection significances of $2.0\sigma$, $1.3\sigma$, and $2.6\sigma$.

For an ideal projection, corresponding to matched true and fiducial 4PCFs, a strong detection significance would be found at small $N_{\rm eig}$ (whereupon the effects of likelihood non-Gaussianity are suppressed), which would increase only slightly as $N_{\rm eig}$ was increased. This behavior is clearly not observed for any of the templates discussed above, suggesting that the signal causing the 4PCF detection is not close to one of the above forms. This is unsurprising: even if the signal is cosmological in nature, it could be sourced by a wide variety of physical effects, each of which has a different signature in the $N_\zeta = 1288$-dimensional 4PCF statistic. That we do not observe a strong detection with any template is additionally consistent with the notion that, if such an effect is real, it lies on the threshold of that detectable by current data.

### E. Tests on mock catalogs

A useful end-to-end test of our analysis pipeline is to apply it to a set of (parity-conserving) mock catalogs and test whether a signal is observed. For this purpose, we will use the NSERIES simulation suite[16]: a set of 84 mock catalogs created to verify the BOSS analysis pipeline [87]. These were constructed from full $N$-body simulations (though are not quite independent) using the cosmological parameters $\{\Omega_m = 0.286, \sigma_8 = 0.82, n_s = 0.97, h = 0.7, \sum m_\nu = 0\}$, and have similar halo occupation distribution and selection function to the BOSS data, as well as a careful treatment of fiber collisions. The NSERIES window function is somewhat different from that of BOSS (cf. Sec. III A), and includes only the NGC region. For this reason, the simulations are accompanied by 2048 PATCHY simulations generated with the NSERIES window function (hereafter "PATCHY-NSERIES" simulations), which will be used to generate covariances and perform the relevant rank tests. Unlike for BOSS, the PATCHY-NSERIES simulations were not calibrated to the small-scale clustering of the NSERIES simulations (which have a slightly modified cosmology to PATCHY-NSERIES); as a result, the covariance of NSERIES and PATCHY-NSERIES could differ somewhat, which would have implications for the parity-violation tests, as discussed in Sec. VI C.

Using the methodology of Sec. IV, we analyze each of the 84 NSERIES simulations in turn, giving the 4PCF measurements shown in Fig. 8 (which display no obvious signal). Here, we show results only for the rank tests; implementation of the other tests discussed above (such as the compressed Gaussian analysis and constraints on the inflationary models of Sec. VII) can be found online.[17] Figure 9(a) compares the distribution of the *pseudo-*$\chi^2$ statistic for the NSERIES and PATCHY-NSERIES simulations (following Sec. V). If our analysis pipeline is working as expected, and the statistical properties of NSERIES and PATCHY-NSERIES are similar, the empirical *pseudo-*$\chi^2$ distributions of the two should match; thus the NSERIES simulations should have a mean rank of $\approx 1024/2048$ (with some deviations expected from sample variance). In practice, we find a mean rank of $244/2048$ (11.9%), which is significantly below the expected value. Since the test statistic is a quadratic form, this does not imply a false detection of parity violation (which would give a mean rank *greater* than $1024/2048$); instead it highlights differences between the NSERIES and PATCHY-NSERIES simulation suites (despite the similarity observed in Fig. 8). As mentioned above, these could come from a variety of factors, such as a different underlying cosmology, the inclusion of fiber collisions in the former, and a different halo occupation distribution. Notably, none of the NSERIES mocks have a mean rank *above* 95%; hence we do not find evidence for parity violation in any case.

To more closely assess the impact of differences in the NSERIES and PATCHY-NSERIES covariance matrices (which cause the difference in *pseudo-*$\chi^2$, in the Gaussian limit), we additionally perform a rank test with the mocks rescaled using the realization-dependent factor of Sec. VI C. The relevant normalization is $0.90 \pm 0.03$ for NSERIES and $0.92 \pm 0.04$ for PATCHY-NSERIES. This differences implies that the NSERIES histogram will shift slightly in the direction of PATCHY-NSERIES; this is observed in Fig. 9(b). The mean rank is $634/2048$ (31.0%), which is significantly closer to the expected value than without rescaling. Around

[16] Available at www.ub.edu/bispectrum/page11.html.

[17] github.com/oliverphilcox/Parity-Odd-4PCF.

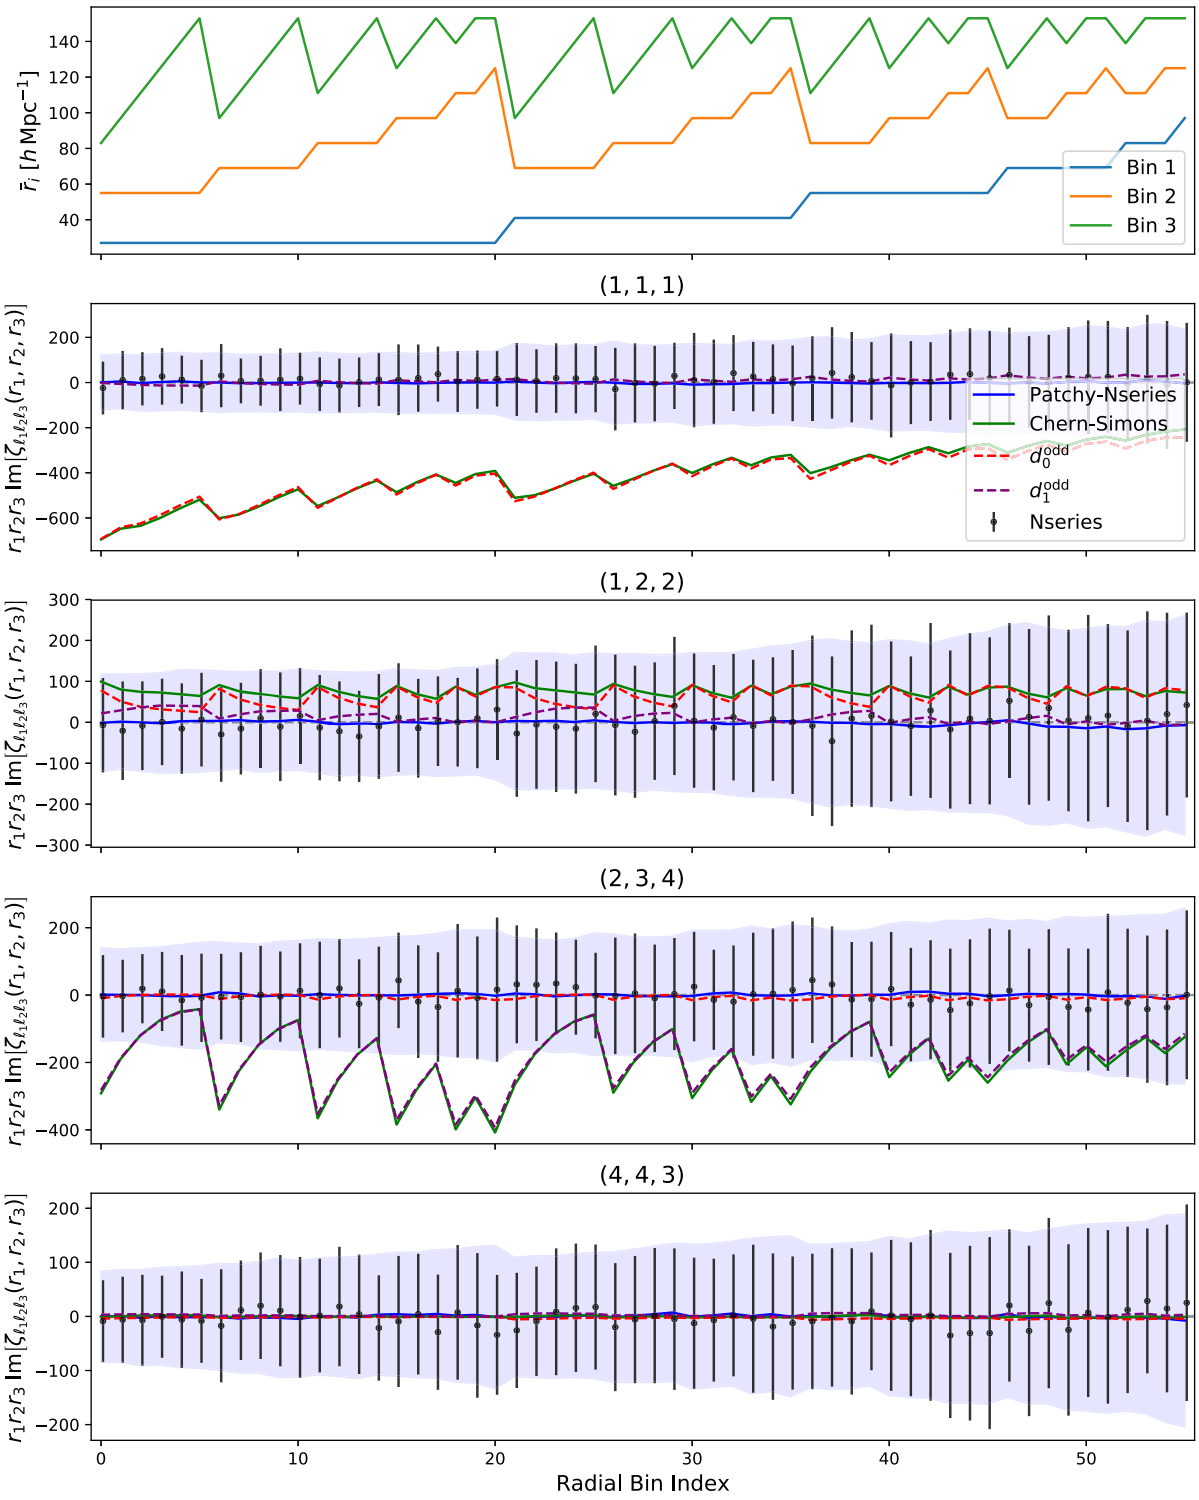


FIG. 8. Measurements of the parity-odd 4PCF from the mean of 84 NSERIES mocks (black) and 2048 NSERIES-PATCHY simulations (blue), following the form of Fig. 2. We additionally display theoretical predictions for the Chern-Simons inflationary model (green) and its two constituent parts, proportional to $d_0^{\rm odd}$ (red dashed lines) and $d_1^{\rm odd}$ (purple dashed lines). Theory models are multiplied by an amplitude corresponding to $A_{\rm CS} = 5 \times 10^4$ for visibility, with data-driven constraints on this parameter presented in Fig. 10. The NSERIES dataset appears consistent with parity conservation, as expected; this is explored further in Fig. 9.

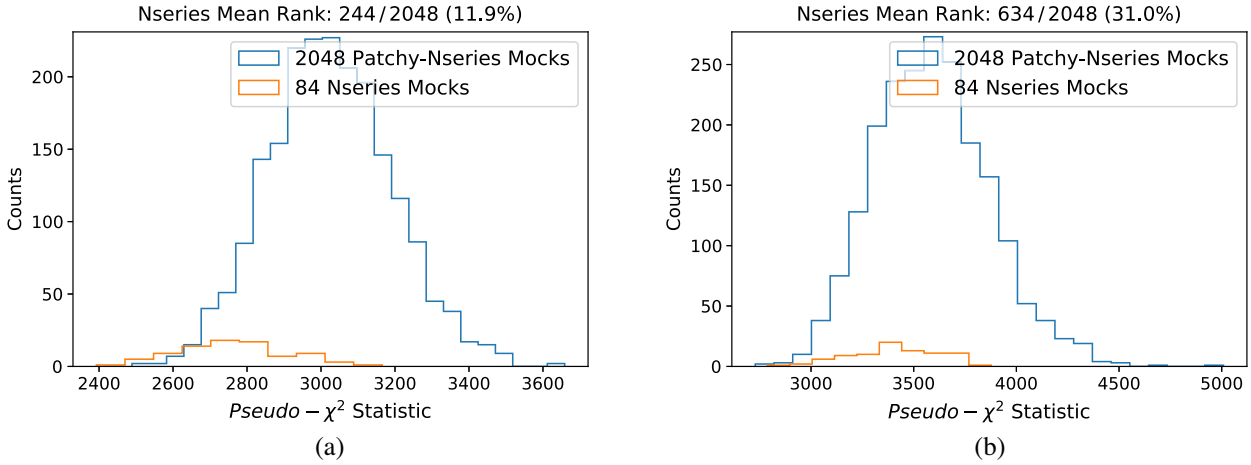


FIG. 9. Distribution of the *pseudo*-$\chi^2$ statistic for a set of high-fidelity NSERIES simulations (orange) and the corresponding empirical distributions obtained from 2048 PATCHY-NSERIES mocks. In the left panel (analogous to Fig. 4), we find a mean rank of 244/2048, with 30/84 mocks lying in the outer 5% region of the empirical distribution PATCHY-NSERIES distribution. As such, we report an *underestimate* of the test statistic relative to that expected. This is attributed to the difference in cosmology and bias parameters between the NSERIES mocks and PATCHY simulations. The right panel shows the effect of including the realization-dependent rescaling factor discussed in Sec. VI C: this reduces the difference between datasets and increases the mean rank to 634/2048, with only 8 of the 84 simulations in the outer 5% of the empirical distribution. In both cases, there are no NSERIES simulations with a *pseudo*-$\chi^2$ value in the upper 5% region of the empirical distribution, and thus no detection of parity violation. (a) Fiducial (b) Including realization-dependent rescaling.

10% of the NSERIES mocks lie in the outer 5% region of the PATCHY-NSERIES histogram, giving approximately a one-in-ten chance of a false detection (or here, an antidetection). The shift in the mean rank induced by the rescaling factor leads to two conclusions: (a) the factor can be usefully adopted to remove the lowest-order differences in simulation covariance matrices (validating the approach of Sec. VI C), and (b) the difference between NSERIES and PATCHY-NSERIES covariances is not fully captured by the rescaling, and is thus scale dependent.

The above tests indicate that the validity of our null tests depend strongly on whether the statistical properties of the simulation suite employed to compute empirical distributions match those of the real data. A difference in the sample covariance matrix (which could be caused by various effects, as elaborated upon below) can lead to a false under- or overdetection of parity violation. For the NSERIES simulation suite, these differences are quite significant and lead to a marked underdetection; however, this can be substantially reduced by including the realization-dependent rescaling factor of Sec. VI C.

### F. Other sources of systematics

While the above tests constrain a variety of different systematic effects, they are by no means exhaustive. Below, we discuss a number of additional effects that could contribute to the tentative detections reported in Sec. V. These can be separated into two groups: (a) effects that lead to the PATCHY simulations having a different covariance (and higher-order statistical properties) to that of the data, and (b) observational phenomena causing a parity-violating signal in the odd-parity 4PCF itself.

First, we consider possible causes for a systematic difference between PATCHY mocks and BOSS data. Besides simple differences in cosmological parameters or halo occupation distributions, discrepancies could arise due to the treatment of nonlinear evolution. The PATCHY mocks are generated using approximate gravity solvers only (before being calibrated to an $N$-body simulation, and the observational data); while this will not affect the large (linear) scales, it can modify the small-scale clustering, and thus provide an error on small scales, where the constraining power of the data is greatest. Simply calibrating the two- and three-point functions does not guarantee that the covariance is well reproduced, since this depends on $N$-point functions up to $N = 8$. Such effects would distort the *pseudo*-$\chi^2$ distribution of the PATCHY mocks and could lead to false detections of parity violation. This hypothesis can be probed via the NSERIES tests of Sec. VI E; once realization-dependent rescaling is included, we find fair agreement between the mock data and PATCHY-NSERIES simulations, and no spurious detections of parity violation. For the BOSS data, we expect better agreement with PATCHY, given that the simulations were calibrated to the observed small-scale clustering; however, we still find a $2.9\sigma$ detection of parity violation, or $2.4\sigma$ when the rescaling factor is included. This reduces the likelihood that the effect is caused by inaccuracies in PATCHY (considering also the scale-cut information discussed in Sec. VI B), though this effect is certainly worthy of future study.

Another effect not included in the PATCHY mocks is that of fiber collisions, arising from the inability to position telescope fibers within a certain distance from each other. If the assignment of fibers to a telescope image is performed in a particular direction (i.e., picking all objects above a certain brightness from one side of a field of view to another), and the quartet of galaxies contains some distant component, a parity-violating signal may be created, due to a preference of one tetrahedral handedness over the other. Two lines of evidence suggest this may not be an important contribution: first, the NSERIES mocks include the effect of fiber collisions and do not show strong evidence for a parity-violating signal; second, most tetrahedra considered in this work have large radial separations. Although fiber collisions happen in the angular domain, rather than radial, all galaxies have minimum separations above 20 $h^{-1}$ Mpc (Sec. VI B), and thus most fibers of relevance will be spaced by tens of arcminutes. The effect could also change the statistical properties of the data (i.e., the covariance matrix); the impact of this is analogous to the effects discussed above.

The BOSS dataset is also known to contain systematic effects on large scales due to imperfectly subtracted foreground modes, arising, for example, from Galactic emission or varying seeing conditions. Given that these do not impact BAO measurements, they are usually ignored, though they may be of more relevance to the analyses considered herein. For such an effect to show up in the odd-parity 4PCF signal, it would need to be parity violating. Since observational effects are not required to obey isotropy and homogeneity, this is possible and could be formed, for example, from the composition of two small-scale 2PCFs (of different lengths) with a large-scale gradient between them. This is analogous to the disconnected 4PCF contribution (but parity odd) and could be sourced by the above observational phenomena, or even cosmological effects such as isocurvature modes. From the analysis employing hemispherical cuts (Sec. VI A), it is clear that, if this was the cause of the parity-breaking detection, it is not a one-off phenomenon, and, moreover, it is not sourced solely by very small or very large tetrahedron configurations (Sec. VI B). The particular form of our tetrahedral basis (i.e., the decomposition into coupled spherical harmonics) makes checking individual tetrahedral configurations difficult, unless a dedicated analysis is performed; the best way to probe them would be with mocks including all observational effects, though none of this type currently exist. We note, however, that such large-scale modes would likely have an impact also on lower-order statistics which utilize large-scale data. The consistency of such two-point and three-point function analyses with those of the CMB suggest that these effects, if present, are comparatively small [76,88].

## VII. CONSTRAINTS ON INFLATIONARY PARITY VIOLATION

The 4PCF measurements presented in Sec. V may be used to place constraints on specific models of cosmological parity violation, such as those involving inflation. As noted in Appendix B of [7], parity violation cannot be generated by single-field inflation,[18] and thus its detection in data could give evidence for multiple fields active in the early universe. Here, we consider a particular multifield model, which gives an analytic form for the parity-odd galaxy 4PCF, in addition to lower-dimensional observables. An analogous procedure may be used to constrain any model which induces a nontrivial parity-odd 4PCF; a selection of these are briefly discussed in Sec. VII A 2.

### A. Primordial correlators

#### 1. Chern-Simons inflation

Consider an inflationary Lagrangian containing the following couplings between an inflaton field, $\phi$, and a $U(1)$ gauge field, $A_\mu$:

$$\mathcal{L} \supset f(\phi)\left[-\frac{1}{4}F_{\mu\nu}F^{\mu\nu} + \frac{\gamma}{4}F_{\mu\nu}\tilde{F}^{\mu\nu}\right], \tag{24}$$

where $F_{\mu\nu} \equiv \partial_\mu A_\nu - \partial_\nu A_\mu$ is a two-form. The second term involves the Hodge dual, $\tilde{F}^{\mu\nu}$, and is of the Chern-Simons form [10,10,13,90,91]. This is controlled by two pieces: a function $f(\phi)$ giving the time evolution of the field, and a dimensionless ratio $\gamma$, which sets the amplitude of parity breaking.[19] Following [13,93], we will assume $\gamma$ to be constant (on naturalness grounds), and fix $f(\phi) \propto a^{-4}$ (for scale factor $a$), such that the vector field has a time-independent vacuum expectation value (VEV) and thus a scale-invariant correlator.[20] If one instead assumes $f(\phi) \propto$ const, the energy density of the vector field will decay as $a^{-4}$ during inflation, giving observational signatures only on the largest scales [12].

The Lagrangian (24) leads to a number of modifications to the standard inflationary picture. First, the presence of a background VEV $A_0^\mu$ (usually represented in the electromagnetic notation as $\boldsymbol{E}_0$, with $\boldsymbol{B}_0 = \boldsymbol{0}$) leads to anisotropic expansion, and its perturbations can provide an isocurvature source for the primordial curvature $\zeta$. Such effects are strongly constrained by CMB data, limiting the energy density in the gauge field (hereafter $\rho_E$) to be a small fraction of the inflaton energy density $\rho_\phi$. Couplings between the inflaton and gauge field will additionally

[18] An exception can occur for ghost inflation [89]; this will be discussed in future work.

[19] Constraints on the reheating temperature from big bang nucleosynthesis restrict the coupling strength to $|\gamma| < 5.5$ [12,92].

[20] This is a natural choice within the Ratra mechanism [93,94].

generate gravitational waves through the metric tensor $h_{ij}$, as well as scalar-tensor couplings. This can lead to observable signatures in CMB $E$-modes and $B$-modes (and nonzero parity-breaking spectra such as $C_\ell^{TB}$ and $C_\ell^{EB}$), though such effects are slow-roll suppressed [13]. We note that a nonzero VEV $\boldsymbol{E}_0$ is a natural prediction of the theory; this is simply the impact of long-wavelength classical perturbations in the vector field which have not yet reentered the horizon.[21] Furthermore, the action of nonzero $\gamma$ is to produce an excess of one gauge field helicity mode over the other, causing a parity asymmetry.

In this work, we are interested in the scalar correlators generated by the above interaction, i.e., only those involving the curvature $\zeta$. As demonstrated in [12,13,38,95], the gauge fields lead to anisotropy in the two-point function:

$$\langle\zeta(\boldsymbol{k}_1)\zeta(\boldsymbol{k}_2)\rangle = (2\pi)^3\delta_{\rm D}(\boldsymbol{k}_1+\boldsymbol{k}_2)P_\zeta(k_1)[1+g_*(\hat{\boldsymbol{k}}_1\cdot\hat{\boldsymbol{E}}_0)], \tag{25}$$

where $P_\zeta \approx H^2/(4\epsilon M_p^2 k^3)$ is the primordial power spectrum, for Hubble parameter $H$, Planck mass $M_p$, and slow-roll parameter $\epsilon$. This is of the well-known Ackerman-Carroll-Wise form [96], for approximately scale-invariant coupling $g_* \propto \rho_E/\rho_\phi$. Furthermore, the Lagrangian given in (24) generates an angle-averaged bispectrum and trispectrum of the curvature perturbation $\zeta$. The first takes the form of [97,98]

$$\begin{aligned}\langle\zeta(\boldsymbol{k}_1)\zeta(\boldsymbol{k}_2)\zeta(\boldsymbol{k}_3)\rangle = (2\pi)^3\delta_{\rm D}(\boldsymbol{k}_1+\boldsymbol{k}_2+\boldsymbol{k}_3)\\ \times\sum_{n=0,1,2} c_n\mathcal{L}_n(\hat{\boldsymbol{k}}_1\cdot\hat{\boldsymbol{k}}_2)\\ \times P_\zeta(k_1)P_\zeta(k_2) + 2\,\text{perms.},\end{aligned} \tag{26}$$

where $\mathcal{L}_n$ is a Legendre polynomial of order $n$ and the coefficients $c_n$ (simply related to *Planck*'s $f_{\rm NL}^{L=n}$ parameters [98]) are again proportional to the fractional energy density in the vector field [13] and satisfy $c_0 = -2/3c_1 = 2c_2$. These parametrize a number of effects beyond the inflationary Lagrangian (24), such as curvature fluctuations sourced by primordial magnetic fields [17], solid inflation [99] (which boasts $c_2 \gg c_0$), and massive spinning particles [24,100]. Finally, the four-point function for the CS model was computed in [38,95] and can be expressed in terms of the reduced trispectrum $t$, defined as

$$\left\langle\prod_{i=1}^4\zeta(\boldsymbol{k}_i)\right\rangle = (2\pi)^3\delta_{\rm D}\left(\sum_{i=1}^4\boldsymbol{k}_i\right)[t^{\boldsymbol{k}_1\boldsymbol{k}_2}_{\boldsymbol{k}_3\boldsymbol{k}_4} + 23\,\text{perms.}]. \tag{27}$$

[21]Following [13], the magnitude of this ought to scale as $N_{\rm tot}-N$, where $N$ is the number of $e$-folds before the CMB modes exited the horizon and $N_{\rm tot}$ gives the total duration of inflation.

Separating even and odd parts, we have

$$\begin{aligned}t^{\boldsymbol{k}_1\boldsymbol{k}_2}_{\boldsymbol{k}_3\boldsymbol{k}_4}\Big|_{\rm even} &= \sum_{n=0,2} d_n^{\rm even}(\mathcal{L}_n(\hat{\boldsymbol{k}}_1\cdot\hat{\boldsymbol{k}}_3)+\mathcal{L}_n(\hat{\boldsymbol{k}}_1\cdot\hat{\boldsymbol{K}})\\ &\quad+\mathcal{L}_n(\hat{\boldsymbol{k}}_3\cdot\hat{\boldsymbol{K}}))P_\zeta(k_1)P_\zeta(k_3)P_\zeta(K),\\ t^{\boldsymbol{k}_1\boldsymbol{k}_2}_{\boldsymbol{k}_3\boldsymbol{k}_4}\Big|_{\rm odd} &= i\sum_{n=0,1} d_n^{\rm odd}(\mathcal{L}_n(\hat{\boldsymbol{k}}_1\cdot\hat{\boldsymbol{k}}_3)+\mathcal{L}_n(\hat{\boldsymbol{k}}_1\cdot\hat{\boldsymbol{K}})\\ &\quad+(-1)^n\mathcal{L}_n(\hat{\boldsymbol{k}}_3\cdot\hat{\boldsymbol{K}}))[(\hat{\boldsymbol{k}}_1\times\hat{\boldsymbol{k}}_3)\cdot\hat{\boldsymbol{K}}]\\ &\quad\times P_\zeta(k_1)P_\zeta(k_3)P_\zeta(K),\end{aligned} \tag{28}$$

where $\boldsymbol{K} = \boldsymbol{k}_3+\boldsymbol{k}_4$ is the trispectrum diagonal and $t^{-\boldsymbol{k}_1-\boldsymbol{k}_2}_{-\boldsymbol{k}_3-\boldsymbol{k}_4} = \pm t^{\boldsymbol{k}_1\boldsymbol{k}_2}_{\boldsymbol{k}_3\boldsymbol{k}_4}$ for the even and odd components, respectively. This depends on new parameters, $\{d_n\}$, with the odd-parity contributions sourced iff $\gamma$ is nonzero (i.e., if the Lagrangian contains a Chern-Simons term). While (28) is a relatively general parametrization for scale-invariant parity-odd primordial trispectra, we may specialize to the Chern-Simons model by fixing $d_0^{\rm even} = d_2^{\rm even}/2$, and $d_0^{\rm odd} = -d_1^{\rm odd}/3$, each of which is linear in $\rho_E/\rho_\phi$.

For later convenience, we will rewrite the odd-parity reduced trispectrum corresponding to (24) as

$$\begin{aligned}t^{\boldsymbol{k}_1\boldsymbol{k}_2}_{\boldsymbol{k}_3\boldsymbol{k}_4}\Big|_{\rm odd} &= 3iA_{\rm CS}P_\zeta(k_1)P_\zeta(k_3)P_\zeta(s)[(\hat{\boldsymbol{k}}_1\times\hat{\boldsymbol{k}}_3)\cdot\hat{\boldsymbol{s}}]\\ &\quad\times[1-\hat{\boldsymbol{k}}_1\cdot\hat{\boldsymbol{k}}_3+\hat{\boldsymbol{k}}_1\cdot\hat{\boldsymbol{s}}-\hat{\boldsymbol{k}}_3\cdot\hat{\boldsymbol{s}}] \equiv t(\boldsymbol{k}_1,\boldsymbol{k}_3,\boldsymbol{s}),\end{aligned} \tag{29}$$

introducing the Mandelstam variable $\boldsymbol{s}\equiv\boldsymbol{k}_1+\boldsymbol{k}_2$, and defining the overall amplitude $A_{\rm CS} \equiv -d_0^{\rm odd} = d_1^{\rm odd}/3$ as

$$\begin{aligned}A_{\rm CS} &\approx \frac{0.3}{\pi^2}\frac{e^{8\pi|\gamma|}}{|\gamma|^6}\frac{|g_*|}{0.01}\left(\frac{N}{60}\right)^2,\\ g_* &\approx -\frac{5.4\times10^5}{\pi}\frac{e^{4\pi|\gamma|}}{|\gamma|^3}\frac{0.01}{\epsilon}\left(\frac{N}{60}\right)^2\frac{\rho_E}{\rho_\phi},\end{aligned} \tag{30}$$

for $|\gamma|>1$, assuming that modes of interest exited the horizon $N$ $e$-folds before the end of inflation. As expected from Sec. II, the reduced trispectrum in (29) is pure imaginary and depends on a parity-odd cross product.

Combining the above, we can constrain the gauge field energy density from the two-, three-, or four-point functions of the scalar field, or, following the transformations outlined in Sec. VII B, the galaxy overdensity. As in [13,38], we expect the constraints to be a strong function of $|\gamma|$: gauge field production increases with $\gamma$, and the $N$-point function scales as $2(N-1)$ powers of the gauge field perturbation $\delta E$, leading to a prefactor of $e^{4\pi(N-1)|\gamma|}/|\gamma|^{3(N-1)}$. In this work, we will derive constraints only from the parity-odd 4PCF: this provides both a tight constraint for $|\gamma|>0$ due to the above argument and is a clean probe, since it is not contaminated by gravitational non-Gaussianity.

#### 2. Other sources

Before proceeding to derive constraints on the Chern-Simons inflationary model, we briefly discuss a number of alternative phenomena that can source parity violation. First, the Lagrangian presented in (24) remains applicable when $\phi$ is *not* the inflaton. This has been suggested as a mechanism for primordial magnetogenesis [93,94] and will source similar parity violation. An additional case of interest is when $F_{\mu\nu}$ is the electromagnetic tensor and $\phi$ an axionlike particle in the late universe; this can generate detectable cosmic birefringence [25,33,101]. Second, the inflationary Lagrangian could contain a *gravitational* Chern-Simons term

$$\mathcal{L} \supset f'(\phi) R^{\lambda}_{\sigma\mu\nu} \tilde{R}^{\sigma\mu\nu}_{\lambda} \tag{31}$$

[9–11], for some Riemann curvature tensor $R_{\mu\nu\rho\sigma}$ and a function $f'(\phi)$ of the inflaton $\phi$. Following a similar calculation to that of [38] for the Lagrangian described above, one can compute the scalar trispectrum corresponding to (31), which is now sourced by couplings to tensor modes rather than vectors.

Primordial parity violation can also arise from *particle exchange* in the early universe (as part of the so-called "cosmological collider" [24,102]). At high energies, it is natural to assume that the inflaton is coupled to additional fields (be they scalars, vectors, or tensors), via some three-point interaction vertices, e.g., $(\partial_\mu\phi)(\partial_\nu\phi)\partial^\mu V^\nu$ for a primordial vector $V^\mu$. In this case, the two-point function of the inflaton (and thus the curvature perturbation $\zeta$) can receive an off-diagonal contribution. A simple example of this occurs for light mediators, which takes the form

$$\langle \zeta(\boldsymbol{k}_1)\zeta(\boldsymbol{k}_2)\rangle'|_{X_p(\boldsymbol{K})} = \frac{B_{\zeta\zeta X_p}(k_1,k_2,K)}{P_{X_p}(K)} X^*_p(\boldsymbol{K})(2\pi)^3\delta_{\rm D}(\boldsymbol{k}_1+\boldsymbol{k}_2+\boldsymbol{K}), \tag{32}$$

where $P_{X_p}$ is the power spectrum of $X$ in some polarization state $p$ and $B_{\zeta\zeta X_p}$ is the interaction three-point function [23,46,56,103,104], or, more precisely, its consistency-relation-violating component [105]. Combining two of these interactions, one obtains an exchange diagram for $\zeta$, i.e., a four-point function of the scalar curvature, with some shape proportional to $|B_{\zeta\zeta X_p}|^2/P_{X_p}$. If the field in question is helical, this can lead to a parity-violating trispectrum if $P_{X_+} \neq P_{X_-}$. As discussed in [23], such contributions can be measured directly using a statistical anisotropy ("fossil") estimator, which can be recast as a model-specific compression of the full four-point function. In the squeezed limit ($K \ll k_1 \approx k_2$, again noting that massive mediators can have different behaviors; cf. [103,104,106]), the interaction bispectrum is approximately given by $(3/2)P_\zeta(k_1)P_{X_p}(K)\epsilon^p_{ij}\hat{k}^i_1\hat{k}^j_1$ [23,24,102], where $\epsilon^p$ is the polarization tensor. For a light vector exchange field $X$, this leads to the following schematic form for the trispectra (following [23]):

$$\begin{aligned} t^{\boldsymbol{k}_1\boldsymbol{k}_2}_{\boldsymbol{k}_3\boldsymbol{k}_4}\Big|_{\rm even} &\sim [P_{X_+}(K)+P_{X_-}(K)]P_\zeta(k_1)P_\zeta(k_3)(\hat{\boldsymbol{k}}_1\cdot\hat{\boldsymbol{K}})(\hat{\boldsymbol{k}}_3\cdot\hat{\boldsymbol{K}})[\hat{\boldsymbol{k}}_1\cdot\hat{\boldsymbol{k}}_3-(\hat{\boldsymbol{k}}_1\cdot\hat{\boldsymbol{K}})(\hat{\boldsymbol{k}}_3\cdot\hat{\boldsymbol{K}})], \\ t^{\boldsymbol{k}_1\boldsymbol{k}_2}_{\boldsymbol{k}_3\boldsymbol{k}_4}\Big|_{\rm odd} &\sim -i[P_{X_+}(K)-P_{X_-}(K)]P_\zeta(k_1)P_\zeta(k_3)(\hat{\boldsymbol{k}}_1\cdot\hat{\boldsymbol{K}})(\hat{\boldsymbol{k}}_3\cdot\hat{\boldsymbol{K}})[(\hat{\boldsymbol{k}}_1\times\hat{\boldsymbol{k}}_3)\cdot\hat{\boldsymbol{K}}]. \end{aligned} \tag{33}$$

The similarities of this and (28) are manifest, particularly if one assumes a scale invariant form for the power spectrum of $X$, such that $P_{X_\pm}(K) \propto P_\zeta(K)$. We thus note that a simple extension to the parametrization of (28) can incorporate trispectra arising from particle exchange. A similar conclusion holds also for intermediate fields $X$ of higher spin (noting that scalar exchange cannot generate parity-violating couplings).

### B. 4PCF model

To place constraints on the Chern-Simons interaction, we must compute the galaxy 4PCF associated with the primordial trispectrum of (29). At redshift $z$, the tree-level galaxy trispectrum, $T_g$, can be related to the primordial correlator via

$$(2\pi)^3\delta_{\rm D}(\boldsymbol{k}_1+\boldsymbol{k}_2+\boldsymbol{k}_3+\boldsymbol{k}_4)T_g(\boldsymbol{k}_1,\boldsymbol{k}_2,\boldsymbol{k}_3,\boldsymbol{k}_4,z) = \left\langle \prod_{i=1}^{4} Z_1(\boldsymbol{k}_i,z)M(k_i,z)\zeta(\boldsymbol{k}_i)\right\rangle, \tag{34}$$

where $M(k,z)$ is the transfer function, defined by $\delta_{\rm matter}(\boldsymbol{k},z) \equiv M(k,z)\zeta(\boldsymbol{k})$, and $Z_1(\boldsymbol{k},z) \equiv b(z)+f(z)(\hat{\boldsymbol{k}}\cdot\hat{\boldsymbol{n}})^2$ is the tree-level galaxy RSD kernel [for linear bias $b(z)$, growth-rate $f(z)$, and line-of-sight $\hat{\boldsymbol{n}}$]. From (34), we may obtain the odd-parity 4PCF using Fourier transforms:

$$\zeta_-(\boldsymbol{r}_1,\boldsymbol{r}_2,\boldsymbol{r}_3,z)=\prod_{i=1}^{4}\left[\int_{\boldsymbol{k}_i}M(k_i,z)Z_1(\boldsymbol{k}_i,z)\right]\int_{\boldsymbol{s}}t(\boldsymbol{k}_1,\boldsymbol{k}_3,\boldsymbol{s})[e^{i(\boldsymbol{k}_1\cdot\boldsymbol{r}_1+\boldsymbol{k}_2\cdot\boldsymbol{r}_2+\boldsymbol{k}_3\cdot\boldsymbol{r}_3)}+23\,\text{perms.}]$$
$$\times(2\pi)^3\delta_{\rm D}(\boldsymbol{k}_1+\boldsymbol{k}_2-\boldsymbol{s})(2\pi)^3\delta_{\rm D}(\boldsymbol{k}_3+\boldsymbol{k}_4+\boldsymbol{s}),\tag{35}$$

shifting the permutation sum to the exponential term by symmetry, and using Dirac delta functions to enforce $\boldsymbol{s}=\boldsymbol{k}_1+\boldsymbol{k}_2=-\boldsymbol{k}_3-\boldsymbol{k}_4$. The corresponding multiplets, $\zeta_{\ell_1\ell_2\ell_3}$, can then be estimated by performing weighted integrals over $\boldsymbol{r}_i$, as in (7). Following a lengthy derivation sketched in Appendix B, we obtain the final form for odd $\ell_1+\ell_2+\ell_3$:

$$\zeta_{\ell_1\ell_2\ell_3}(r_1,r_2,r_3,z)=(4\pi)^{11}\sqrt{2}A_{\rm CS}i^{\ell_1+\ell_2+\ell_3}\sum_{L_1L_2L_3L_4L_5L_5'}i^{L_1+L_2+L_3+L_4-L_5+L_5'}\mathcal{C}_{L_1L_2L_3L_4L_5L_5'}$$
$$\times\begin{pmatrix}L_1&L_2&L_5\\0&0&0\end{pmatrix}\begin{pmatrix}L_3&L_4&L_5'\\0&0&0\end{pmatrix}\mathcal{M}^{\ell_1\ell_2(\ell_3)\ell_30}_{L_1L_2L_3L_4L_5L_5'}$$
$$\times\int x^2dx\int x'^2dx'K_{L_5L_5'}(x,x')I^{\ell_1}_{L_1}(x;r_1)J^{\ell_2}_{L_2}(x;r_2)I^{\ell_3}_{L_3}(x';r_3)J^{0}_{L_4}(x';0)+23\,\text{perms.}\tag{36}$$

[cf. (B9)]. Here, $I$, $J$, and $K$ are Bessel-weighted integrals over the transfer function and/or primordial power spectrum (B10), $\mathcal{C}_{\ell_1\cdots\ell_N}\equiv\sqrt{(2\ell_1+1)\cdots(2\ell_N+1)}$, and $\mathcal{M}$ is a coupling matrix given in (B11). The 4PCF may thus be computed as a two-dimensional integral, following evaluation of the (radially binned) $I$, $J$, and $K$ functions for a range of values of $L$, $\ell$, $x$, and $x'$. In practice, the 4PCF model is computed in PYTHON, with the various Wigner $3-j$ and $9-j$ symbols evaluated using SYMPY.[22] In Fig. 8 we plot the theoretical model for a range of multiplets, finding a shape that depends strongly on both $\{r_i\}$ and $\{\ell_i\}$, with some multiplets dominated by the $d_0^{\rm odd}$ part, and others by that proportional to $d_1^{\rm odd}$.

### C. Amplitude constraints

The importance of the inflationary gauge field may be quantified by the ratio of energy densities, $\rho_E/\rho_\phi$, or the parity-odd 4PCF amplitude, $A_{\rm CS}$ (30). This sets the level of parity violation imprinted in the primordial inflaton perturbations, arising from interactions with the $U(1)$ gauge field. To constrain the amplitudes, we perform parameter inference using the measured 4PCF multiplets of Sec. V and the Chern-Simons model given in (36). For simplicity, we will assume the data to be Gaussian distributed and work in a compressed subspace containing $N_{\rm eig}=100$ basis vectors for each of the NGC and SGC regions (whose distribution was shown to be approximately Gaussian in Sec. V). In Sec. IV B, a minimum-variance criterion was used to select the basis vectors; here, we instead pick those with maximal signal to noise for the Chern-Simons model. The reduced dimensionality facilitates direct use of the PATCHY simulations to form the sample covariance (20); to account for the finite number of mocks, we perform inference using the following log-likelihood:

$$-\log L(A_{\rm CS})=\frac{N_{\rm mocks}}{2}\log\left[1+\frac{T^2(A_{\rm CS})}{N_{\rm mocks}-1}\right]+\text{const.}\tag{37}$$

This uses the $T^2$ statistic, defined analogously to (21):

$$T^2(A_{\rm CS})\equiv(\hat{v}_{\rm data}-A_{\rm CS}v_{\rm CS})^T\hat{C}_v^{-1}(\hat{v}_{\rm data}-A_{\rm CS}v_{\rm CS}),\tag{38}$$

where $\hat{v}_{\rm data}$ represents the compressed 4PCF data vector, $\hat{C}_v$ is a sample covariance, and $v_{\rm CS}$ is the compressed 4PCF model of (36), excluding the $A_{\rm CS}$ prefactor. Likelihoods for the NGC and SGC regions are constructed separately and multiplied, assuming independence. Here, we perform two analyses: one using the BOSS data, and the second using the mean of $N_{\rm mocks}=2048$ PATCHY mocks. In the latter case, no Chern-Simons contribution should be present.

Figure 10 shows the resulting constraints on the trispectrum amplitude. For the mean of PATCHY mocks, the $1\sigma$ constraint is $A_{\rm CS}=0\pm760$ (demonstrating unbiasedness, as expected), with $A_{\rm CS}=570\pm780$ observed for BOSS. In both cases, the constraints are consistent with zero, suggesting that the Chern-Simons coupling is not responsible for the detection of parity violation reported in Sec. V. If we additionally restrict to $A_{\rm CS}\geq0$, we find that $A_{\rm CS}<1500$ and $A_{\rm CS}<2000$ for the mean-of-PATCHY and BOSS datasets, respectively (95% C.L.).[23] Additionally, including tetrahedra with small separations between secondary galaxies (as in Sec. V) does not appreciably improve the

[22]Code implementing this calculation is available at github.com/oliverphilcox/Parity-Odd-4PCF.

[23]We have additionally verified that no false detection of Chern-Simons inflation is obtained when using the NSERIES mocks of Sec. VI E.

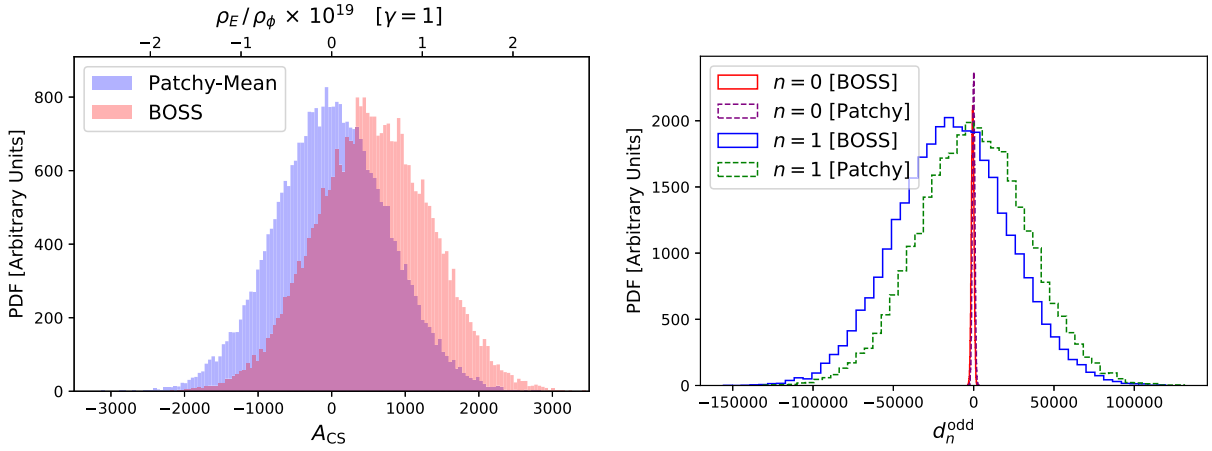


FIG. 10. Constraints on the amplitude of physical models for parity violation, for both the Chern-Simons Lagrangian of (24) (left) and the more general parametrization of (28) (right), using the explicit parity-odd 4PCF prediction given in (36) and plotted in Fig. 8. To obtain these distributions, we fit the measured 4PCF multiplets shown in Fig. 2 to templates derived in Appendix B, following compression of both observations and model into a 100-dimensional subspace. For the Chern-Simons model, we give results for both the trispectrum amplitude $A_{\rm CS}$ and the corresponding ratio of parity-breaking gauge field and inflaton energy densities, $\rho_E/\rho_\phi$, using relation (30) with $|\gamma| = 1$. The $1\sigma$ constraints are $A_{\rm CS} = 0 \pm 760$ for the mean of 2048 PATCHY mocks (blue) and $A_{\rm CS} = 570 \pm 780$ for the BOSS data (red), both of which are consistent with zero. The right panel gives constraints on the $d_n^{\rm odd}$ parameters appearing in (28): we find $d_0^{\rm odd} = -610 \pm 800$ ($-20 \pm 770$) and $d_1^{\rm odd} = -13\,000 \pm 35\,000$ ($0 \pm 34\,000$) for BOSS (mean of Patchy) data, respectively. In the Chern-Simons model, $-d_0^{\rm odd} = d_1^{\rm odd}/3 = A_{\rm CS}$.

constraints, which we justify by noting that the bulk of the signal to noise occurs on comparatively large scales.

Of greater physical interest are the constraints on the energy densities $\rho_E/\rho_\phi$. These may be obtained from the $A_{\rm CS}$ bounds using (30), assuming fiducial values for the inflationary parameters and fixing the coupling strength to $\gamma = 1$, giving an equal contribution from the parity-even and parity-odd terms in (24).[24] Using the BOSS CMASS data, we find $\rho_E/\rho_\phi < 1.6 \times 10^{-19}$ (95% C.L.). If $|\gamma|$ is increased to 2, the gauge field production is strongly amplified, and the constraint tightens to $\rho_E/\rho_\phi < 3.5 \times 10^{-33}$.

We may additionally place limits on the phenomenological parameters $\{d_n^{\rm odd}\}$ appearing in (28). Using an analogous method to the above, we find the $1\sigma$ constraints $d_0^{\rm odd} = -610 \pm 800$ and $d_1^{\rm odd} = -13\,000 \pm 35\,000$ from BOSS, both of which are fully consistent with zero. For the mean of 2048 Patchy simulations, we find $d_0^{\rm odd} = -20 \pm 770$ and $d_1^{\rm odd} = 0 \pm 34\,000$, again indicating that the method is unbiased. Although the physical scale of the two coefficients is the same (in the Chern-Simons model they follow the relation $-d_0^{\rm odd} = d_1^{\rm odd}/3 = A_{\rm CS}$), the first parameter is constrained almost 50× better than the second: this is attributed to the different angular behavior of the two terms, with $d_0^{\rm odd}$ dominating the $\{\ell_1, \ell_2, \ell_3\} = \{1, 1, 1\}$ multiplet, for example (cf. Fig. 8).

Our results for the Chern-Simons model may be compared to those obtained from the power spectrum and bispectrum of the CMB.[25] In particular, the *Planck* 2018 $T$- and $E$-mode datasets (analyzed with the SMICA prescription) gave the constraints

$$-0.036 \le g_* \le 0.036, \qquad -13 \le c_0 \le 11,$$
$$-7 \le c_1 \le 281, \qquad -55 \le c_2 \le 37 \tag{39}$$

[98,107], on the inflationary parameters $g_*$ and $\{c_n\}$ appearing in the two- and three-point parametrizations of (25) and (26) at 95% C.L., and translating into our notation. Assuming $\gamma = 1$ and the above fiducial parameters, these can be used to place bounds on the gauge field energy density: $\rho_E/\rho_\phi \lesssim 7 \times 10^{-13}$ ($2 \times 10^{-16}$) using the two- (three-)point function measurements (cf. [38]). Furthermore, a forecast of the detectability of $\rho_E/\rho_\phi$ from the CMB four-point function was presented in [38]. This predicted a bound on $\rho_E/\rho_\phi$ of $\sim 3 \times 10^{-20}$ at 95% C.L. [or equivalently $\sigma(d_1^{\rm odd}) = 640$] for a cosmic-variance dominated measurement using $\ell_{\rm max} = 2000$ and $|\gamma| = 1$.

Bounds on the gauge field energy density from the BOSS 4PCF are far stronger than those obtained from the CMB anisotropic power spectrum and isotropic bispectrum (for $|\gamma| > 0$), due to the exponential dependence on $|\gamma|$ [[38], Fig. 4]. As such, they represent the strongest current

[24]Note that exchanging $\gamma \to -\gamma$ simply swaps the dominant and suppressed helicity states of the gauge field.

[25]Since the coupling is assumed to be active only during inflation, the Lagrangian (24) does not generate cosmic birefringence. Generation of parity-violating CMB spectra is possible, however (due to helical gravitational wave production), but this is slow-roll suppressed, as noted above.

constraints on Chern-Simons inflationary models. Our measurement is roughly a factor of 5 worse than that predicted for the CMB: this occurs since the latter contains significantly more Fourier modes than the observed galaxy distribution, and thus an increased signal-to-noise ratio (although it is subject to the smoothing effects of projection integrals). As the volume of spectroscopic data grows, we expect the constraints on $\rho_E/\rho_\phi$ to significantly strengthen, especially considering that the signal to noise of the Chern-Simons 4PCF (36) scales as $\sqrt{V_{\rm survey}}$, roughly independent of redshift.[26] A survey such as DESI will probe $\sim$100$\times$ the BOSS volume [47] and should thus be expected to tighten the constraints on $\rho_E/\rho_\phi$ (and any other parity-breaking model amplitudes) by roughly an order of magnitude, providing stronger constraints on parity-breaking inflation than possible with the CMB. Finally, we note that, for LSS, the parity-odd 4PCF is an optimal place in which to search for these signatures, since, unlike other observables, the statistic is free from gravitational effects; thus, we do not have to marginalize over the effects of late-time non-Gaussianity.

## VIII. SUMMARY AND CONCLUSIONS

Searching for parity violation provides a unique manner in which to probe new physics occurring in the early universe, including that of multifield inflation, baryogenesis, and primordial magnetic field generation. While there is a long history of constraining various parity-breaking phenomena using the CMB [6,9,10,13,17,20,37,38,108], few analyses have made use of LSS data. In this work, we have placed the first constraints on (hitherto unexplored) scalar-type parity violation using the BOSS CMASS galaxy sample. The isotropic NPCFs are parity sensitive only if $N > 3$ [23,38,39]; recent developments in NPCF computation [51] have enabled efficient computation of the galaxy 4PCF, enabling such analyses. To provide a model-agnostic test, we have performed a blind search for parity violation using the full parity-odd 4PCF (whose expectation value is zero in $\Lambda$CDM). Our primary tool has been a nonparametric rank test, comparing the BOSS 4PCF (on scales between 20 $h^{-1}$ Mpc and 160 $h^{-1}$ Mpc) to that of a suite of realistic mock catalogs. This avoids the need to assume a Gaussian likelihood, and provides a robust (albeit conservative) model-agnostic test. In the BOSS sample we found tentative evidence for parity violation, with a detection probability of 99.6%, equivalent to $2.9\sigma$. This indicates either new physics beyond the standard model or unknown systematics.

As an additional test, we have performed a classical $\chi^2$-based analysis of the BOSS data, making use of a data compression scheme and a covariance matrix computed from mock catalogs. Furthermore, we use a theoretical Gaussian covariance [57] to facilitate high-fidelity compression, which reduces the dataset to $N_{\rm eig}$ numbers; importantly, the results are *not* biased for this choice, avoiding a potential systematic error. Postcompression, the empirical $\chi^2$ distribution closely matches that of the theory model for $N_{\rm eig} \lesssim 100$; this gives credence to the assumption of Gaussianity. For this test, we find a detection probability of 83.3% from BOSS when using $N_{\rm eig} = 100$, or 100.0% when using $N_{\rm eig} = 250$, though the latter may be artificially inflated from likelihood non-Gaussianity. The results are broadly consistent with those from the rank test; however, the strong dependence on $N_{\rm eig}$ implies that our basis decomposition is inefficient, and that information may be being lost.

We have carried out a number of tests to explore potential systematic effects in our data which could lead to a false detection of parity violation. These include splitting the data into subregions, imposing radial and angular cuts, comparing against mock catalogs, altering the compression scheme, and normalizing by an overall rescaling factor. No clear evidence for systematics is observed, and we find our detection to be relatively coherent across various scales and sky regions. That said, our tests do rely heavily on the PATCHY mocks well representing the statistical properties of the BOSS data; although we have ruled out differences due to an overall rescaling factor, a scale dependent difference remains the most likely cause of our results, in the absence of a cosmological signal.

Finally, we have used the measured 4PCF to bound the amplitudes of physical models of parity violation. Here, we have primarily considered a single scenario; a Chern-Simons term in the inflationary Lagrangian, which couples the inflaton to a $U(1)$ gauge field. This leads to a definite prediction for the primordial polyspectra [13,38], which, with some effort, can be translated into a model for the galaxy 4PCF. Performing a Gaussian likelihood analysis using this template gives a comparable constraint on the ratio of gauge field and inflaton energy densities to that expected from the CMB [38] (but much stronger than that from lower-order statistics); $\rho_E/\rho_\phi < 1.6 \times 10^{-19}$ (95% C.L.), assuming standard inflationary parameters. Notably, this does not appear to explain the above parity-violating signal. Similar constraints may be obtained for any other physical model giving a definite prediction for the galaxy 4PCF.

The coming years will lead to an explosion in the volume of LSS data available, which will either confirm or refute the tentative detection of parity violation found herein. Unlike for the gravitational contribution [54], the signal-to-noise of the inflationary 4PCF is not a strong function of redshift, with the constraints on models of new physics being primarily sensitive to the survey volume. To further increase the constraining power, we may fold in additional information, for example, using the 5PCF and anisotropic NPCFs (which source additional information regarding vector parity breaking [43]). Going beyond spectroscopic

[26]Note that this differs from the signal to noise of the *gravitational* 4PCF, which scales as $[b(z)D(z)]^2\sqrt{V_{\rm survey}}$ [54].

surveys, it is likely that high-volume observables such as intensity mapping and the Lyman-$\alpha$ forest, as well as the CMB itself, will shed additional light on this, pinning down a variety of new physics models.

## ACKNOWLEDGMENTS

We thank Jiamin Hou for providing us with analytic covariance matrices. We are additionally indebted to Stephon Alexander, Giovanni Cabass, He Jia, Marc Kamionkowski, Morgane König, and David Spergel for insightful discussions regarding parity violation; Robert Cahn, Daniel Eisenstein, Jiamin Hou, and Zachary Slepian for conversations relating to NPCFs; William Underwood for suggesting the nonparametric statistics of Sec. IV A; and Mikhail Ivanov, Marko Simonović, and Matias Zaldarriaga for general feedback. Furthermore, we are grateful to Paolo Campeti, Johannes Eskilt, and the anonymous referee for providing valuable feedback on the manuscript. O. H. E. P. acknowledges funding from the WFIRST program through NNG26PJ30C and NNN12AA01C, thanks the Simons Foundation for additional support, and is grateful to the Institute for Advanced Study for their hospitality and abundance of baked goods. The authors are pleased to acknowledge that the work reported in this paper 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. Funding for SDSS-III has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation, and the U.S. Department of Energy Office of Science. The SDSS-III web site is [113]. SDSS-III is managed by the Astrophysical Research Consortium for the Participating Institutions of the SDSS-III Collaboration including the University of Arizona, the Brazilian Participation Group, Brookhaven National Laboratory, Carnegie Mellon University, University of Florida, the French Participation Group, the German Participation Group, Harvard University, the Instituto de Astrofisica de Canarias, the Michigan State/Notre Dame/JINA Participation Group, Johns Hopkins University, Lawrence Berkeley National Laboratory, Max Planck Institute for Astrophysics, Max Planck Institute for Extraterrestrial Physics, New Mexico State University, New York University, Ohio State University, Pennsylvania State University, University of Portsmouth, Princeton University, the Spanish Participation Group, University of Tokyo, University of Utah, Vanderbilt University, University of Virginia, University of Washington, and Yale University. The massive production of all MultiDark-Patchy mocks for the BOSS Final Data Release has been performed at the BSC Marenostrum supercomputer, the Hydra cluster at the Instituto de Fısica Teorica UAM/CSIC, and NERSC at the Lawrence Berkeley National Laboratory. We acknowledge support from the Spanish MICINNs Consolider-Ingenio 2010 Programme under Grant No. MultiDark CSD2009-00064, MINECO Centro de Excelencia Severo Ochoa Programme under Grants No. SEV- 2012-0249, and No. AYA2014-60641-C2-1-P. The MultiDark-Patchy mocks was an effort led from the IFT UAM-CSIC by F. Prada's group (C.-H. Chuang, S. Rodriguez-Torres, and C. Scoccola) in collaboration with C. Zhao (Tsinghua U.), F.-S. Kitaura (AIP), A. Klypin (NMSU), G. Yepes (UAM), and the BOSS galaxy clustering working group.

## APPENDIX A: LIKELIHOOD NON-GAUSSIANITY

A major assumption of most cosmological analyses is that the underlying likelihood for the statistic of interest can be well approximated as Gaussian. While this is often ensured by the central limit theorem, it can break down in the case of highly correlated data, such as that considered in this work. In this appendix, we present a simple test to check whether the likelihood of the full 4PCF can be justifiably considered Gaussian.

For this purpose, we take the PATCHY-NSERIES NGC simulations (Sec. VI E) and partition them into two groups. The first 500 are used to compute a sample mean and covariance for the distribution, thus defining a Gaussian distribution from which we draw $10^5$ mock observations. Each of these is compressed into one dimension via the *pseudo-*$\chi^2$ statistic defined in (18) and histogrammed. This can then be compared to the empirical *pseudo-*$\chi^2$ distribution obtained from the remaining simulations directly. If the likelihood is Gaussian, the two distributions should match.

The resulting PDFs are shown in Fig. 11. Notably, the empirical and Gaussianized distributions do not match, indicating that the full 4PCF distribution is not well described by a Gaussian, even in the best-case scenario

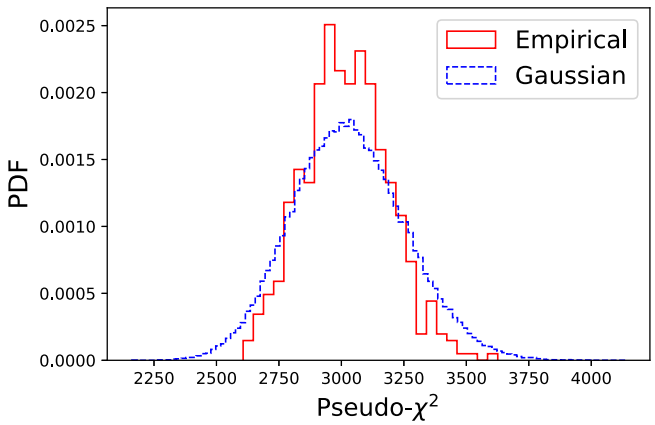


FIG. 11. Distribution of the *pseudo-*$\chi^2$ statistic (18) estimated from PATCHY-NSERIES simulations (red) and Gaussian realizations (blue), drawn from a distribution with mean and covariance estimated from a second set of PATCHY-NSERIES simulations. The clear differences between the distributions indicates that the full 4PCF likelihood cannot be well-approximated as Gaussian.

when the mean and variance are estimated from the simulations. In particular, the Gaussian assumption overestimates the sample variance, which will reduce any potential detection significance. If one instead uses the *theoretical* covariance to define the Gaussian distribution, the situation is far worse: the *pseudo-*$\chi^2$ distribution instead peaks at $\approx$1250, indicating a breakdown of the modeling assumptions (as discussed in Sec. III C). Finally, we note that, although the full (uncompressed) 4PCF distribution appears to be non-Gaussian, this does not imply that the same is true for the projected statistics of Sec. IV B, since the central limit theorem becomes more applicable as the dimensionality reduces.

## APPENDIX B: DERIVATION OF THE CHERN-SIMONS 4PCF MODEL

Below, we sketch the derivation of the parity-odd 4PCF induced by the Chern-Simons coupling of Sec. VII. Our starting point is the general expression given in (35), which is a product of four pieces. By expanding the angular dependence of each term using the isotropic basis functions (Sec. II A), we may compute the full 4PCF efficiently.

To begin, we consider the primordial Chern-Simons trispectrum defined in (29). The angular pieces may be written in terms of isotropic basis functions of three coordinates using [52] [Appendix A.2]:

$$(\hat{\boldsymbol{k}}_1\times\hat{\boldsymbol{k}}_3)\cdot\hat{\boldsymbol{s}} = i\frac{\sqrt{2}}{3}(4\pi)^{3/2}\mathcal{P}_{111}(\hat{\boldsymbol{k}}_1,\hat{\boldsymbol{k}}_3,\hat{\boldsymbol{s}}), \qquad \hat{\boldsymbol{k}}_1\cdot\hat{\boldsymbol{k}}_3 = -\frac{1}{\sqrt{3}}(4\pi)^{3/2}\mathcal{P}_{110}(\hat{\boldsymbol{k}}_1,\hat{\boldsymbol{k}}_3,\hat{\boldsymbol{s}}),$$
$$\hat{\boldsymbol{k}}_1\cdot\hat{\boldsymbol{s}} = -\frac{1}{\sqrt{3}}(4\pi)^{3/2}\mathcal{P}_{101}(\hat{\boldsymbol{k}}_1,\hat{\boldsymbol{k}}_3,\hat{\boldsymbol{s}}), \qquad \hat{\boldsymbol{k}}_1\cdot\hat{\boldsymbol{k}}_3 = -\frac{1}{\sqrt{3}}(4\pi)^{3/2}\mathcal{P}_{011}(\hat{\boldsymbol{k}}_1,\hat{\boldsymbol{k}}_3,\hat{\boldsymbol{s}}). \tag{B1}$$

The resulting products of two basis functions can be simplified using [ [52], Sec. VI C], yielding

$$t(\boldsymbol{k}_1,\boldsymbol{k}_3,\boldsymbol{s}) = -\sqrt{2}A_{\rm CS}(4\pi)^{3/2}P_\zeta(k_1)P_\zeta(k_3)P_\zeta(s) \times\left[\mathcal{P}_{111}(\hat{\boldsymbol{k}}_1,\hat{\boldsymbol{k}}_3,\hat{\boldsymbol{s}}) + \frac{1}{\sqrt{5}}\mathcal{P}_{221}(\hat{\boldsymbol{k}}_1,\hat{\boldsymbol{k}}_3,\hat{\boldsymbol{s}}) + \frac{1}{\sqrt{5}}\mathcal{P}_{212}(\hat{\boldsymbol{k}}_1,\hat{\boldsymbol{k}}_3,\hat{\boldsymbol{s}}) - \frac{1}{\sqrt{5}}\mathcal{P}_{122}(\hat{\boldsymbol{k}}_1,\hat{\boldsymbol{k}}_3,\hat{\boldsymbol{s}})\right], \tag{B2}$$

whose radial part is separable in $k_1$, $k_3$, and $s$. For the general parity-odd trispectrum given in (28), we obtain the same result, but with the replacement $A_{\rm CS}\mathcal{P}_{111}\to -d_0^{\rm odd}\mathcal{P}_{111}$, $A_{\rm CS}\mathcal{P}_{221}+2\,\text{perms.}\to d_1^{\rm odd}\mathcal{P}_{221}/3+2\,\text{perms.}$

Next, we consider the Dirac delta functions. By rewriting $(2\pi)^3\delta_{\rm D}(\boldsymbol{k}_1+\boldsymbol{k}_2-\boldsymbol{s})$ as a complex exponential, inserting plane-wave expansions [ [109], Eq. 16.63], then performing the angular integral, the function can be expressed as a sum over one-dimensional integrals and isotropic basis functions of three coordinates:

$$(2\pi)^3\delta_{\rm D}(\boldsymbol{k}_1+\boldsymbol{k}_2-\boldsymbol{s}) = (4\pi)^{5/2}\sum_{L_1L_2L_5} i^{L_1+L_2+L_5}(-1)^{L_5}\mathcal{C}_{L_1L_2L_5}\begin{pmatrix}L_1 & L_2 & L_5\\ 0 & 0 & 0\end{pmatrix}\mathcal{P}_{L_1L_2L_5}(\hat{\boldsymbol{k}}_1,\hat{\boldsymbol{k}}_2,\hat{\boldsymbol{s}}) \times\int_0^\infty x^2dx\,j_{L_1}(k_1x)j_{L_2}(k_2x)j_{L_5}(sx), \tag{B3}$$

where $\mathcal{C}_{L_1\cdots L_N}\equiv\sqrt{(2L_1+1)\cdots(2L_N+1)}$, the $2\times 3$ matrices are Wigner 3-$j$ symbols, and we have used properties of the Gaunt integral [ [110], Eq. 34.3.22]. Similarly,

$$(2\pi)^3\delta_{\rm D}(\boldsymbol{k}_3+\boldsymbol{k}_4+\boldsymbol{s}) = (4\pi)^{5/2}\sum_{L_3L_4L_5'} i^{L_3+L_4+L_5'}\mathcal{C}_{L_3L_4L_5'}\begin{pmatrix}L_3 & L_4 & L_5'\\ 0 & 0 & 0\end{pmatrix}\mathcal{P}_{L_3L_4L_5'}(\hat{\boldsymbol{k}}_3,\hat{\boldsymbol{k}}_4,\hat{\boldsymbol{s}}) \times\int_0^\infty x'^2dx'\,j_{L_3}(k_3x')j_{L_4}(k_4x')j_{L_5'}(sx'). \tag{B4}$$

Note that the integrands are again separable in $\{k_i\}$ and $s$. Using the approach of [111], they may equivalently be rewritten as infinite sums.

For the transfer functions $M(k,z)Z_1(\boldsymbol{k},z)$, we first expand the redshift-space kernel $Z_1(\boldsymbol{k})$ in spherical harmonics (dropping the redshift dependence for clarity):

$$Z_1(\hat{\boldsymbol{k}};\hat{\boldsymbol{n}}) \equiv b + f(\hat{\boldsymbol{k}}\cdot\hat{\boldsymbol{n}})^2 = 4\pi\sum_{\ell m}\left[\left(b+\frac{f}{3}\right)\delta^{\rm K}_{\ell 0} + \frac{2f}{15}\delta^{\rm K}_{\ell 2}\right]Y^*_{\ell m}(\hat{\boldsymbol{n}})Y_{\ell m}(\hat{\boldsymbol{k}}) \equiv 4\pi\sum_{\ell m} Z_\ell Y^*_{\ell m}(\hat{\boldsymbol{n}})Y_{\ell m}(\hat{\boldsymbol{k}}), \tag{B5}$$

for linear bias $b(z)$, growth rate $f(z)$, and line-of-sight $\hat{\boldsymbol{n}}$. Since we consider only isotropic 4PCFs in this work, we can integrate over the Line-of-sight orientation (which is equivalent to performing a joint rotation of all $\{\boldsymbol{r}_i\}$). Following some algebra, this leads to a set of isotropic functions of *four* coordinates (see [52] for details):

$$\begin{aligned}\int\frac{d\hat{\boldsymbol{n}}}{4\pi}Z_1(\hat{\boldsymbol{k}}_1;\hat{\boldsymbol{n}})Z_1(\hat{\boldsymbol{k}}_2;\hat{\boldsymbol{n}})Z_1(\hat{\boldsymbol{k}}_3;\hat{\boldsymbol{n}})Z_1(\hat{\boldsymbol{k}}_4;\hat{\boldsymbol{n}}) &= (4\pi)^2\sum_{j_1j_2j_{12}j_3j_4}\begin{pmatrix} j_1 & j_2 & j_{12}\\ 0 & 0 & 0\end{pmatrix}\begin{pmatrix} j_{12} & j_3 & j_4\\ 0 & 0 & 0\end{pmatrix}Z_{j_1}Z_{j_2}Z_{j_3}Z_{j_4}\\ &\times \mathcal{C}_{j_1j_2j_{12}j_3j_4}\mathcal{P}_{j_1j_2(j_{12})j_3j_4}(\hat{\boldsymbol{k}}_1,\hat{\boldsymbol{k}}_2,\hat{\boldsymbol{k}}_4,\hat{\boldsymbol{k}}_4),\end{aligned} \tag{B6}$$

where $j_i \in \{0,2\}$ and $j_{12}\in\{0,2,4\}$.

The final contribution is from the Fourier basis functions and their permutations, which can be written

$$e^{i(\boldsymbol{k}_1\cdot\boldsymbol{r}_1+\boldsymbol{k}_2\cdot\boldsymbol{r}_2+\boldsymbol{k}_3\cdot\boldsymbol{r}_3)} + 23\,\text{perms.} = \sum_H e^{i(\boldsymbol{k}_1\cdot\boldsymbol{r}_{H1}+\boldsymbol{k}_2\cdot\boldsymbol{r}_{H2}+\boldsymbol{k}_3\cdot\boldsymbol{r}_{H3}+\boldsymbol{k}_4\cdot\boldsymbol{r}_{H4})}, \tag{B7}$$

where $\{H1,H2,H3,H4\}$ is one of the 24 permutations of $\{1,2,3,4\}$, and we have introduced $\boldsymbol{r}_4 = 0$ for convenience. Projecting onto the 4PCF basis functions $\mathcal{P}_{\ell_1\ell_2\ell_3}(\hat{\boldsymbol{r}}_1,\hat{\boldsymbol{r}}_2,\hat{\boldsymbol{r}}_3)$ gives a sum of isotropic functions of four coordinates:

$$\begin{aligned}&\sum_H(4\pi)^{-1/2}\int d\hat{\boldsymbol{r}}_1 d\hat{\boldsymbol{r}}_2 d\hat{\boldsymbol{r}}_3 d\hat{\boldsymbol{r}}_4\mathcal{P}^*_{\ell_1\ell_2\ell_3}(\hat{\boldsymbol{r}}_1,\hat{\boldsymbol{r}}_2,\hat{\boldsymbol{r}}_3)e^{i(\boldsymbol{k}_1\cdot\boldsymbol{r}_{H1}+\boldsymbol{k}_2\cdot\boldsymbol{r}_{H2}+\boldsymbol{k}_3\cdot\boldsymbol{r}_{H3}+\boldsymbol{k}_4\cdot\boldsymbol{r}_{H4})}Y_{\ell_4m_4}(\hat{\boldsymbol{r}}_4)\\ &\quad=\sum_H(4\pi)^{7/2}\Phi_H(-i)^{\ell_1+\ell_2+\ell_3}j_{\ell_{H1}}(k_1r_{H1})j_{\ell_{H2}}(k_2r_{H2})j_{\ell_{H3}}(k_3r_{H3})j_{\ell_{H4}}(k_4r_{H4})\mathcal{P}_{\ell_{H1}\ell_{H2}(\ell^*)\ell_{H3}\ell_{H4}}(\hat{\boldsymbol{k}}_1,\hat{\boldsymbol{k}}_2,\hat{\boldsymbol{k}}_3,\hat{\boldsymbol{k}}_4),\end{aligned} \tag{B8}$$

using the plane wave expansion and inserting $\ell_4 = m_4 = r_4 = 0$ in the first line. In the second line we include a permutation factor $\Phi_H$, given by $(-1)^{\ell_1+\ell_2+\ell_3}$ if $\{\ell_{H1},\ell_{H2},\ell_{H3},\ell_{H4}\}$ is an odd permutation of $\{\ell_1,\ell_2,\ell_3\}$ (removing the zero element) and unity else. Furthermore, $\ell^*$ is set by the position of the zero, e.g., $\ell^* = \ell_{H2}$ if $\ell_{H1} = 0$, $\ell^* = \ell_{H4}$ if $\ell_{H3} = 0$, etc.

Combining the above results, we obtain

$$\begin{aligned}\zeta_{\ell_1\ell_2\ell_3}(r_1,r_2,r_3) &= (4\pi)^{11}\sqrt{2}A_{\rm CS}i^{\ell_1+\ell_2+\ell_3}\sum_H\sum_{L_1L_2L_3L_4L_5L_5'}i^{L_1+L_2+L_3+L_4-L_5+L_5'}\mathcal{C}_{L_1L_2L_3L_4L_5L_5'}\\ &\times\begin{pmatrix} L_1 & L_2 & L_5\\ 0 & 0 & 0\end{pmatrix}\begin{pmatrix} L_3 & L_4 & L_5'\\ 0 & 0 & 0\end{pmatrix}\times\mathcal{M}^{\ell_{H1}\ell_{H2}(\ell^*)\ell_{H3}\ell_{H4}}_{L_1L_2L_3L_4L_5L_5'}\\ &\times\int x^2dx\int x'^2dx' K_{L_5L_5'}(x,x')I^{\ell_{H1}}_{L_1}(x;r_{H1})J^{\ell_{H2}}_{L_2}(x;r_{H2})I^{\ell_{H3}}_{L_3}(x';r_{H3})J^{\ell_{H4}}_{L_4}(x';r_{H4}),\end{aligned} \tag{B9}$$

defining the integrals

$$\begin{aligned}I^\ell_L(x;r) &\equiv \int_0^\infty\frac{k^2dk}{2\pi^2}M(k)P_\zeta(k)j_L(kx)j_\ell(kr), \qquad J^\ell_L(x;r)\equiv\int_0^\infty\frac{k^2dk}{2\pi^2}M(k)j_L(kx)j_\ell(kr),\\ K_{LL'}(x,x') &\equiv \int_0^\infty\frac{s^2ds}{2\pi^2}P_\zeta(s)j_L(sx)j_{L'}(sx').\end{aligned} \tag{B10}$$

In practice, we must integrate the statistic over radial bins of finite width, which corresponds to replacing, e.g., $j_\ell(kr)$ with $\overline{j}^b_\ell(k)$ for bin $b$. The bin-integrated Bessel functions are analytic and their forms can be found in [112] [Appendix A].

The coupling matrix in (B9) is given by an integral over five isotropic basis functions of five coordinates:

$$\mathcal{M}^{\ell_{H1}\ell_{H2}(\ell^*)\ell_{H3}\ell_{H4}}_{L_1L_2L_3L_4L_5L_5'} = \sum_{j_1j_2j_{12}j_3j_4} \mathcal{C}_{j_1j_2j_{12}j_3j_4} \begin{pmatrix} j_1 & j_2 & j_{12} \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} j_{12} & j_3 & j_4 \\ 0 & 0 & 0 \end{pmatrix} Z_{j_1} Z_{j_2} Z_{j_3} Z_{j_4} \int d\hat{\boldsymbol{k}}_1 d\hat{\boldsymbol{k}}_2 d\hat{\boldsymbol{k}}_3 d\hat{\boldsymbol{k}}_4 d\hat{\boldsymbol{s}}$$
$$\times [\mathcal{P}_{j_1j_2(j_{12})j_3(j_4)j_40} \mathcal{P}_{L_1L_2(L_5)0(L_5)0L_5} \mathcal{P}_{00(0)L_3(L_3)L_4L_5'} \mathcal{P}_{\ell_{H1}\ell_{H2}(\ell^*)\ell_{H3}(\ell_{H4})\ell_{H4}0}](\hat{\boldsymbol{k}}_1, \hat{\boldsymbol{k}}_2, \hat{\boldsymbol{k}}_3, \hat{\boldsymbol{k}}_4, \hat{\boldsymbol{s}})$$
$$\times \left[\mathcal{P}_{10(1)1(1)01} + \frac{1}{\sqrt{5}}\mathcal{P}_{20(2)2(2)01} + \frac{1}{\sqrt{5}}\mathcal{P}_{20(2)1(2)02} - \frac{1}{\sqrt{5}}\mathcal{P}_{10(1)2(2)02}\right](\hat{\boldsymbol{k}}_1, \hat{\boldsymbol{k}}_2, \hat{\boldsymbol{k}}_3, \hat{\boldsymbol{k}}_4, \hat{\boldsymbol{s}}), \tag{B11}$$

where we have noted that isotropic functions of $N$ coordinates may be rewritten in terms of those with $N + M \geq N$ coordinates by inserting a factor $(4\pi)^{M/2}$. Despite its complexity, this can be evaluated analytically, making extensive use of the product relation for isotropic basis functions of five coordinates:

$$\mathcal{P}_\Lambda \mathcal{P}_{\Lambda'} = (4\pi)^{-5/2} \sum_{\Lambda''} (-1)^{\Lambda_1''+\Lambda_2''+\Lambda_3''+\Lambda_4''+\Lambda_5''} \mathcal{C}_\Lambda \mathcal{C}_{\Lambda'} \mathcal{C}_{\Lambda''} \prod_{i=1}^{5} \left[\begin{pmatrix} \Lambda_i & \Lambda_i' & \Lambda_i'' \\ 0 & 0 & 0 \end{pmatrix}\right] \mathcal{P}_{\Lambda''}$$
$$\times \begin{Bmatrix} \Lambda_1 & \Lambda_2 & \Lambda_{12} \\ \Lambda_i' & \Lambda_2' & \Lambda_{12}' \\ \Lambda_1'' & \Lambda_2'' & \Lambda_{12}'' \end{Bmatrix} \begin{Bmatrix} \Lambda_{12} & \Lambda_3 & \Lambda_{123} \\ \Lambda_{12}' & \Lambda_3' & \Lambda_{123}' \\ \Lambda_{12}'' & \Lambda_3'' & \Lambda_{123}'' \end{Bmatrix} \begin{Bmatrix} \Lambda_{123} & \Lambda_4 & \Lambda_5 \\ \Lambda_{123}' & \Lambda_4' & \Lambda_5' \\ \Lambda_{123}'' & \Lambda_4'' & \Lambda_5'' \end{Bmatrix} \tag{B12}$$

[[52], Sec. VI E], where $\Lambda \equiv \{\Lambda_1, \Lambda_2, (\Lambda_3), \Lambda_4, (\Lambda_{123}), \Lambda_4, \Lambda_5\}$, the curly braces indicate Wigner $9-j$ symbols, and $\mathcal{C}_\Lambda$ involves all elements of $\Lambda$. This simplifies considerably when some elements of $\Lambda$ or $\Lambda'$ are zero.


[1] T. D. Lee and C. N. Yang, Question of parity conservation in weak interactions, Phys. Rev. **104**, 254 (1956).
[2] A. D. Sakharov, Violation of *CP* invariance, C asymmetry, and baryon asymmetry of the universe, Sov. J. Exp. Theor. Phys. Lett. **5**, 24 (1967).
[3] S. M. Carroll, Quintessence and the Rest of the World: Suppressing Long-Range Interactions, Phys. Rev. Lett. **81**, 3067 (1998).
[4] H. Davoudiasl, R. Kitano, G. D. Kribs, H. Murayama, and P. J. Steinhardt, Gravitational Baryogenesis, Phys. Rev. Lett. **93**, 201301 (2004).
[5] S. Alexander, Inflationary birefringence and baryogenesis, Int. J. Mod. Phys. D **25**, 1640013 (2016).
[6] S. H. S. Alexander, Is cosmic parity violation responsible for the anomalies in the WMAP data?, Phys. Lett. B **660**, 444 (2008).
[7] T. Liu, X. Tong, Y. Wang, and Z.-Z. Xianyu, Probing P and *CP* violations on the cosmological collider, J. High Energy Phys. 04 (2020) 189.
[8] S. H. Alexander, M. E. Peskin, and M. M. Sheikh-Jabbari, Leptogenesis from Gravity Waves in Models of Inflation, Phys. Rev. Lett. **96**, 081301 (2006).
[9] A. Lue, L. Wang, and M. Kamionkowski, Cosmological Signature of New Parity-Violating Interactions, Phys. Rev. Lett. **83**, 1506 (1999).
[10] N. Bartolo and G. Orlando, Parity breaking signatures from a Chern-Simons coupling during inflation: The case of non-Gaussian gravitational waves, J. Cosmol. Astropart. Phys. 07 (2017) 034.
[11] S. Alexander and N. Yunes, Chern-Simons modified general relativity, Phys. Rep. **480**, 1 (2009).
[12] N. Bartolo, S. Matarrese, M. Peloso, and M. Shiraishi, Parity-violating and anisotropic correlations in pseudoscalar inflation, J. Cosmol. Astropart. Phys. 01 (2015) 027.
[13] N. Bartolo, S. Matarrese, M. Peloso, and M. Shiraishi, Parity-violating CMB correlators with non-decaying statistical anisotropy, J. Cosmol. Astropart. Phys. 07 (2015) 039.
[14] L. Bordin and G. Cabass, Graviton non-Gaussianities and parity violation in the EFT of inflation, J. Cosmol. Astropart. Phys. 07 (2020) 014.
[15] D. Grasso and H. R. Rubinstein, Magnetic fields in the early Universe, Phys. Rep. **348**, 163 (2001).
[16] P. A. R. Ade, N. Aghanim, M. Arnaud, F. Arroja, M. Ashdown *et al.* (Planck Collaboration), Planck 2015 results. XIX. Constraints on primordial magnetic fields, Astron. Astrophys. **594**, A19 (2016).
[17] M. Shiraishi, Parity violation of primordial magnetic fields in the CMB bispectrum, J. Cosmol. Astropart. Phys. 06 (2012) 015.
[18] L. Pogosian and M. Wyman, B-modes from cosmic strings, Phys. Rev. D **77**, 083509 (2008).
[19] D. Regan and M. Hindmarsh, The bispectrum of matter perturbations from cosmic strings, J. Cosmol. Astropart. Phys. 03 (2015) 008.

[20] I. Y. Rybak and L. Sousa, CMB anisotropies generated by cosmic string loops, Phys. Rev. D **104**, 023507 (2021).

[21] J. L. Cook and L. Sorbo, Particle production during inflation and gravitational waves detectable by ground-based interferometers, Phys. Rev. D **85**, 023534 (2012).

[22] P. Adshead, J. Giblin, T. John, T. R. Scully, and E. I. Sfakianakis, Gauge-preheating and the end of axion inflation, J. Cosmol. Astropart. Phys. 12 (2015) 034.

[23] D. Jeong and M. Kamionkowski, Clustering Fossils from the Early Universe, Phys. Rev. Lett. **108**, 251301 (2012).

[24] N. Arkani-Hamed and J. Maldacena, Cosmological collider physics, arXiv:1503.08043.

[25] Y. Minami and E. Komatsu, New Extraction of the Cosmic Birefringence from the Planck 2018 Polarization Data, Phys. Rev. Lett. **125**, 221301 (2020).

[26] J. R. Eskilt and E. Komatsu, Improved constraints on cosmic bBirefringence from the WMAP and Planck cosmic microwave background polarization data, arXiv:2205.13962.

[27] S. E. Clark, C.-G. Kim, J. C. Hill, and B. S. Hensley, The origin of parity violation in polarized dust emission and implications for cosmic birefringence, Astrophys. J. **919**, 53 (2021).

[28] J. R. Eskilt, Frequency-dependent constraints on cosmic birefringence from the LFI and HFI Planck data release 4, Astron. Astrophys. **662**, A10 (2022).

[29] P. Diego-Palazuelos *et al.*, Cosmic Birefringence from the Planck Data Release 4, Phys. Rev. Lett. **128**, 091302 (2022).

[30] U. Seljak and M. Zaldarriaga, Signature of Gravity Waves in the Polarization of the Microwave Background, Phys. Rev. Lett. **78**, 2054 (1997).

[31] R. G. Crittenden, P. Natarajan, U.-L. Pen, and T. Theuns, Discriminating weak lensing from intrinsic spin correlations using the curl-gradient decomposition, Astrophys. J. **568**, 20 (2002).

[32] S. M. Carroll, G. B. Field, and R. Jackiw, Limits on a Lorentz- and parity-violating modification of electrodynamics, Phys. Rev. D **41**, 1231 (1990).

[33] V. Gluscevic and M. Kamionkowski, Testing parity-violating mechanisms with cosmic microwave background experiments, Phys. Rev. D **81**, 123529 (2010).

[34] R. R. Caldwell and C. Devulder, Axion gauge field inflation and gravitational leptogenesis: A lower bound on B modes from the matter-antimatter asymmetry of the Universe, Phys. Rev. D **97**, 023532 (2018).

[35] K. W. Masui, U.-L. Pen, and N. Turok, Two- and Three-Dimensional Probes of Parity in Primordial Gravity Waves, Phys. Rev. Lett. **118**, 221301 (2017).

[36] F. Schmidt, N. E. Chisari, and C. Dvorkin, Imprint of inflation on galaxy shape correlations, J. Cosmol. Astropart. Phys. 2015 (2015) 032.

[37] M. Kamionkowski and T. Souradeep, Odd-parity cosmic microwave background bispectrum, Phys. Rev. D **83**, 027301 (2011).

[38] M. Shiraishi, Parity violation in the CMB trispectrum from the scalar sector, Phys. Rev. D **94**, 083503 (2016).

[39] R. N. Cahn, Z. Slepian, and J. Hou, A test for cosmological parity violation using the 3D distribution of galaxies, arXiv:2110.12004.

[40] N. Kaiser, Clustering in real space and in redshift space, Mon. Not. R. Astron. Soc. **227**, 1 (1987).

[41] A. J. S. Hamilton, Linear Redshift Distortions: A Review, in *The Evolving Universe*, edited by D. Hamilton, Astrophysics and Space Science Library Vol. 231 (Springer, US, 1998), p. 185.

[42] C. Bonvin, R. Durrer, N. Khosravi, M. Kunz, and I. Sawicki, Redshift-space distortions from vector perturbations, J. Cosmol. Astropart. Phys. 02 (2018) 028.

[43] D. Jeong and F. Schmidt, The odd-parity galaxy bispectrum, Phys. Rev. D **102**, 023530 (2020).

[44] D. Munshi, A. Heavens, A. Cooray, J. Smidt, P. Coles, and P. Serra, New optimized estimators for the primordial trispectrum, Mon. Not. R. Astron. Soc. **412**, 1993 (2011).

[45] K. M. Smith, L. Senatore, and M. Zaldarriaga, Optimal analysis of the CMB trispectrum, arXiv:1502.00635.

[46] L. Dai, D. Jeong, and M. Kamionkowski, Seeking inflation fossils in the cosmic microwave background, Phys. Rev. D **87**, 103006 (2013).

[47] A. Aghamousa, J. Aguilar, S. Ahlen, S. Alam, L. E. Allen *et al.* (DESI Collaboration), The DESI experiment part I: Science, targeting, and survey design, arXiv:1611.00036.

[48] R. Laureijs, J. Amiaux, S. Arduini, J. L. Augueres, J. Brinchmann, R. Cole *et al.*, Euclid definition study report, arXiv:1110.3193.

[49] P. A. Abell, J. Allison, S. F. Anderson, J. R. Andrew, J. R. P. Angel *et al.* (LSST Science Collaboration), Lsst science book, version 2.0, 2009.

[50] C. G. Sabiu, B. Hoyle, J. Kim, and X.-D. Li, Graph database solution for higher-order spatial statistics in the era of big data, Astrophys. J. Suppl. Ser. **242**, 29 (2019).

[51] O. H. E. Philcox, Z. Slepian, J. Hou, C. Warner, R. N. Cahn, and D. J. Eisenstein, ENCORE: Estimating galaxy $N$-point correlation functions in $\mathcal{O}(N_g^2)$ time, Mon. Not. R. Astron. Soc. **509**, 2457 (2021).

[52] R. N. Cahn and Z. Slepian, Isotropic N-point basis functions and their properties, arXiv:2010.14418.

[53] O. H. E. Philcox and Z. Slepian, Efficient computation of $N$-point correlation functions in $D$ dimensions, arXiv:2106.10278.

[54] O. H. E. Philcox, J. Hou, and Z. Slepian, A first detection of the connected 4-point correlation function of galaxies using the BOSS CMASS sample, arXiv:2108.01670.

[55] M. Liguori, E. Sefusatti, J. R. Fergusson, and E. P. S. Shellard, Primordial non-Gaussianity and bispectrum measurements in the cosmic microwave background and large-scale structure, Adv. Astron. **2010**, 980523 (2010).

[56] E. Dimastrogiovanni, M. Fasiello, D. Jeong, and M. Kamionkowski, Inflationary tensor fossils in large-scale structure, J. Cosmol. Astropart. Phys. 12 (2014) 050.

[57] J. Hou, R. Cahn, O. Philcox, and Z. Slepian, Analytic gaussian covariance matrices for galaxy $N$-point correlation functions, 10.1103/PhysRevD.106.043515.

[58] T. L. Smith and M. Kamionkowski, Probability distribution for non-Gaussianity estimators constructed from the CMB trispectrum, Phys. Rev. D **86**, 063009 (2012).

[59] C. Hahn, F. Beutler, M. Sinha, A. Berlind, S. Ho, and D. W. Hogg, Likelihood non-Gaussianity in large-scale structure analyses, Mon. Not. R. Astron. Soc. **485**, 2956 (2019).

[60] D. A. Varshalovich, A. N. Moskalev, and V. K. Khersonskii, *Quantum Theory of Angular Momentum* (World Scientific, Singapore, 1988).
[61] S. D. Landy and A. S. Szalay, Bias and variance of angular correlation functions, Astrophys. J. **412**, 64 (1993).
[62] I. Szapudi and A. S. Szalay, A new class of estimators for the N-point correlations, Astrophys. J. Lett. **494**, L41 (1998).
[63] Z. Slepian, D. J. Eisenstein, J. R. Brownstein, C.-H. Chuang, H. Gil-Marín, S. Ho *et al.*, Detection of baryon acoustic oscillation features in the large-scale three-point correlation function of SDSS BOSS DR12 CMASS galaxies, Mon. Not. R. Astron. Soc. **469**, 1738 (2017).
[64] S. Alam, F. D. Albareti, C. Allende Prieto, F. Anders, S. F. Anderson, T. Anderton *et al.*, The eleventh and twelfth data releases of the Sloan Digital Sky Survey: Final data from SDSS-III, Astrophys. J. Suppl. Ser. **219**, 12 (2015).
[65] D. J. Eisenstein, D. H. Weinberg, E. Agol, H. Aihara, C. Allende Prieto, S. F. Anderson *et al.*, SDSS-III: Massive spectroscopic surveys of the distant universe, the milky way, and extra-solar planetary systems, Astron. J. **142**, 72 (2011).
[66] K. S. Dawson, D. J. Schlegel, C. P. Ahn, S. F. Anderson, É. Aubourg, S. Bailey *et al.*, The Baryon oscillation spectroscopic survey of SDSS-III, Astron. J. **145**, 10 (2013).
[67] F. Beutler, H.-J. Seo, S. Saito, C.-H. Chuang, A. J. Cuesta, D. J. Eisenstein *et al.*, The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: Anisotropic galaxy clustering in Fourier space, Mon. Not. R. Astron. Soc. **466**, 2242 (2017).
[68] H. A. Feldman, N. Kaiser, and J. A. Peacock, Power-Spectrum Analysis of Three-dimensional Redshift Surveys, Astrophys. J. **426**, 23 (1994).
[69] F.-S. Kitaura, S. Rodríguez-Torres, C.-H. Chuang, C. Zhao, F. Prada, H. Gil-Marín *et al.*, The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: Mock galaxy catalogues for the BOSS Final Data Release, Mon. Not. R. Astron. Soc. **456**, 4156 (2016).
[70] S. A. Rodríguez-Torres, C.-H. Chuang, F. Prada, H. Guo, A. Klypin, P. Behroozi *et al.*, The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: Modelling the clustering and halo occupation distribution of BOSS CMASS galaxies in the Final Data Release, Mon. Not. R. Astron. Soc. **460**, 1173 (2016).
[71] Z. Slepian and D. J. Eisenstein, Computing the three-point correlation function of galaxies in $O(N\hat{2})$ time, Mon. Not. R. Astron. Soc. **454**, 4142 (2015).
[72] F. Sosa Nuñez and G. Niz, On the fast random sampling and other properties of the three point correlation function in galaxy surveys, J. Cosmol. Astropart. Phys. 12 (2020) 021.
[73] M. Vargas-Magaña, S. Ho, A. J. Cuesta, R. O'Connell, A. J. Ross, D. J. Eisenstein *et al.*, The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: Theoretical systematics and Baryon Acoustic Oscillations in the galaxy correlation function, Mon. Not. R. Astron. Soc. **477**, 1153 (2018).
[74] D. Wadekar and R. Scoccimarro, Galaxy power spectrum multipoles covariance in perturbation theory, Phys. Rev. D **102**, 123517 (2020).
[75] O. H. E. Philcox and D. J. Eisenstein, Estimating covariance matrices for two- and three-point correlation function moments in Arbitrary Survey Geometries, Mon. Not. R. Astron. Soc. **490**, 5931 (2019).
[76] M. M. Ivanov, M. Simonović, and M. Zaldarriaga, Cosmological parameters from the BOSS galaxy power spectrum, J. Cosmol. Astropart. Phys. 05 (2020) 042.
[77] A. Chudaykin, M. M. Ivanov, O. H. E. Philcox, and M. Simonović, Nonlinear perturbation theory extension of the Boltzmann code CLASS, Phys. Rev. D **102**, 063533 (2020).
[78] A. F. Heavens, R. Jimenez, and O. Lahav, Massive lossless data compression and multiple parameter estimation from galaxy spectra, Mon. Not. R. Astron. Soc. **317**, 965 (2000).
[79] R. Scoccimarro, The bispectrum: From theory to observations, Astrophys. J. **544**, 597 (2000).
[80] J. Alsing and B. Wandelt, Generalized massive optimal data compression, Mon. Not. R. Astron. Soc. **476**, L60 (2018).
[81] O. H. E. Philcox, M. M. Ivanov, M. Zaldarriaga, M. Simonović, and M. Schmittfull, Fewer mocks and less noise: Reducing the dimensionality of cosmological observables with subspace projections, Phys. Rev. D **103**, 043508 (2021).
[82] J. Wishart and M. S. Bartlett, The generalised product moment distribution in a normal system, Proc. Cambridge Philos. Soc. **29**, 260 (1933).
[83] J. Hartlap, P. Simon, and P. Schneider, Why your model parameter confidences might be too optimistic. Unbiased estimation of the inverse covariance matrix, Astron. Astrophys. **464**, 399 (2007).
[84] E. Sellentin and A. F. Heavens, Parameter inference with estimated covariance matrices, Mon. Not. R. Astron. Soc. **456**, L132 (2016).
[85] C. Eckart and G. Young, The approximation of one matrix by another of lower rank, Psychometrika **1**, 211 (1936).
[86] O. H. E. Philcox, J. Hou, and Z. Slepian, A first detection of the connected 4-point correlation function of galaxies using the BOSS CMASS sample, arXiv:2108.01670.
[87] S. Alam, M. Ata, S. Bailey, F. Beutler, D. Bizyaev, J. A. Blazek *et al.*, The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: Cosmological analysis of the DR12 galaxy sample, Mon. Not. R. Astron. Soc. **470**, 2617 (2017).
[88] O. H. E. Philcox and M. M. Ivanov, The BOSS DR12 full-shape cosmology: ΛCDM constraints from the large-scale galaxy power spectrum and bispectrum monopole, Phys. Rev. D **105**, 043517 (2022).
[89] N. Arkani-Hamed, P. Creminelli, S. Mukohyama, and M. Zaldarriaga, Ghost inflation, J. Cosmol. Astropart. Phys. 04 (2004) 001.
[90] K. Freese, J. A. Frieman, and A. V. Olinto, Natural inflation with pseudo Nambu-Goldstone bosons, Phys. Rev. Lett. **65**, 3233 (1990).
[91] S. Alexander, A. Marcianò, and D. Spergel, Chern-Simons inflation and baryogenesis, J. Cosmol. Astropart. Phys. 04 (2013) 046.
[92] S. Hannestad, What is the lowest possible reheating temperature?, Phys. Rev. D **70**, 043506 (2004).

[93] C. Caprini and L. Sorbo, Adding helicity to inflationary magnetogenesis, J. Cosmol. Astropart. Phys. 10 (2014) 056.
[94] B. Ratra, Cosmological "Seed" magnetic field from inflation, Astrophys. J. Lett. **391**, L1 (1992).
[95] M. Shiraishi, E. Komatsu, and M. Peloso, Signatures of anisotropic sources in the trispectrum of the cosmic microwave background, J. Cosmol. Astropart. Phys. 04 (2014) 027.
[96] L. Ackerman, S. M. Carroll, and M. B. Wise, Imprints of a primordial preferred direction on the microwave background, Phys. Rev. D **75**, 083502 (2007).
[97] M. Shiraishi, E. Komatsu, M. Peloso, and N. Barnaby, Signatures of anisotropic sources in the squeezed-limit bispectrum of the cosmic microwave background, J. Cosmol. Astropart. Phys. 2013 (2013) 002.
[98] Y. Akrami, F. Arroja, M. Ashdown, J. Aumont, C. Baccigalupi *et al.* (Planck Collaboration), Planck 2018 results. IX. Constraints on primordial non-Gaussianity, Astron. Astrophys. **641**, A9 (2020).
[99] S. Endlich, A. Nicolis, and J. Wang, Solid inflation, J. Cosmol. Astropart. Phys. 2013 (2013) 011.
[100] G. Franciolini, A. Kehagias, A. Riotto, and M. Shiraishi, Detecting higher spin fields through statistical anisotropy in the CMB bispectrum, Phys. Rev. D **98**, 043533 (2018).
[101] P. Campeti, O. Özsoy, I. Obata, and M. Shiraishi, New constraints on axion-gauge field dynamics during inflation from *Planck* and BICEP/Keck data sets, arXiv: 2203.03401.
[102] J. Maldacena, Non-Gaussian features of primordial fluctuations in single field inflationary models, J. High Energy Phys. 05 (2003) 013.
[103] E. Dimastrogiovanni, M. Fasiello, and L. Pinol, Primordial stochastic gravitational wave background anisotropies: In-in formalization and applications, arXiv:2203.17192.
[104] E. Dimastrogiovanni, M. Fasiello, and M. Kamionkowski, Imprints of massive primordial fields on large-scale structure, J. Cosmol. Astropart. Phys. 02 (2016) 017.
[105] P. Creminelli and M. Zaldarriaga, Single field consistency relation for the 3-point function, J. Cosmol. Astropart. Phys. 10 (2004) 006.
[106] D. Bertolini, K. Schutz, M. P. Solon, and K. M. Zurek, The trispectrum in the effective field theory of large scale structure, J. Cosmol. Astropart. Phys. 06 (2016) 052.
[107] Y. Akrami, F. Arroja, M. Ashdown, J. Aumont, C. Baccigalupi *et al.* (Planck Collaboration), Planck 2018 results. X. Constraints on inflation, Astron. Astrophys. **641**, A10 (2020).
[108] L. Sorbo, Parity violation in the cosmic microwave background from a pseudoscalar inflaton, J. Cosmol. Astropart. Phys. 06 (2011) 003.
[109] G. Arfken, H. Weber, and F. Harris, *Mathematical Methods for Physicists: A Comprehensive Guide* (Elsevier Science, 2013).
[110] NIST, NIST Digital Library of Mathematical Functions, DLMF.
[111] O. H. E. Philcox and Z. Slepian, An exact integral-to-sum relation for products of Bessel functions, Proc. R. Soc. A **477**, 20210376 (2021).
[112] O. H. E. Philcox and D. J. Eisenstein, Computing the small-scale galaxy power spectrum and bispectrum in configuration space, Mon. Not. R. Astron. Soc. **492**, 1214 (2020).
[113] www.sdss3.org.

*Correction:* An important referencing omission has been rectified in the previously published last two sentences (now three sentences) of the fifth paragraph in Sec. I.