

Testing graviton parity and Gaussianity with Planck T -, E -, and B -mode bispectra

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

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

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

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

 (Received 5 January 2024; accepted 22 February 2024; published 14 March 2024)

Many inflationary theories predict a non-Gaussian spectrum of primordial tensor perturbations, sourced from nonstandard vacuum fluctuations, modified general relativity, or new particles such as gauge fields. Several such models also predict a chiral spectrum in which one polarization state dominates. In this work, we place constraints on the non-Gaussianity and parity properties of primordial gravitational waves utilizing the *Planck* PR4 temperature and polarization data set. Using recently developed quasioptimal bispectrum estimators, we compute binned parity-even and parity-odd bispectra for all combinations of cosmic microwave background (CMB) T -, E -, and B -modes with $2 \leq \ell < 500$, and perform both blind tests, sensitive to arbitrary three-point functions, and targeted analyses of a well-motivated equilateral gravitational-wave template (sourced by gauge fields) with amplitude $f_{\text{NL}}^{\text{eff}}$. This is the first time B -modes have been included in primordial non-Gaussianity analyses; they are found to strengthen constraints on the parity-even sector by $\simeq 30\%$ and dominate the parity-odd bounds, without inducing bias. We report no detection of non-Gaussianity (of either parity), with the template amplitude constrained to $f_{\text{NL}}^{\text{eff}} = 900 \pm 700$ (stable with respect to a number of analysis variations), compared to 1300 ± 1200 in *Planck* 2018. The methods applied herein can be reapplied to upcoming CMB data sets such as LiteBIRD, with the inclusion of B -modes poised to dramatically improve future bounds on tensor non-Gaussianity.

DOI: [10.1103/PhysRevD.109.063522](https://doi.org/10.1103/PhysRevD.109.063522)

I. INTRODUCTION

High-precision measurements of the cosmic microwave background (CMB) have provided fruitful information on primordial scalar fluctuations, supporting both the Λ CDM model and an inflationary epoch in the very early Universe. In contrast, experiments have provided little insight into the primordial tensor sector due to the nondetection of early-Universe gravitational waves (GWs), despite numerous searches [1,2] (see [3] for a review).

Primordial GWs have been widely regarded as a smoking gun of cosmic inflation since they are naturally generated by quantum fluctuations of the inflationary vacuum. However, they may also be sourced by other mechanisms. For instance, if gauge fields exist and do not decay during inflation, GWs can be sourced by gravitational interactions (e.g., [4,5]). Their production becomes more efficient if the gauge fields couple to axions [5–11] (see also [12,13] for reviews). Furthermore, GWs could be a signature of, e.g., nonstandard primordial vacuum states (such as α vacua) [14,15], nonattractor phases in inflation [16], additional

scalar-tensor and derivative couplings [17–20], or Chern-Simons modified gravity [6,21–24] (see [25] and references therein for other scenarios). Many such contributions can be efficiently modeled in a UV-physics-agnostic manner with an effective field theory of inflation, often invoking bootstrap methods [26–30]. Being able to distinguish between these contributions and those intrinsic to a simple Bunch-Davies vacuum is thus an important issue.

Under standard inflationary assumptions, GWs sourced by the vacuum have a statistically monotonic spectrum, which is both parity symmetric and (approximately) Gaussian [31]. More complex models involving, for example, axion-gauge field interactions can source non-Gaussian and/or parity-violating distributions, whose particular form encodes the microphysics of the early Universe [5–11,14,16,24,32–35] (for example, breaking the tensor consistency relation [36]). Diagnosing the statistical properties of the tensor sector is thus a powerful way to probe the GW origins and shed light on inflationary physics; furthermore, we note that large GW bispectra are not *a priori* ruled out by the current bounds on tensor power spectra (e.g., [10,11,14,16,23,26–29,35,37]). Motivated by this, here we perform an analysis to jointly test the parity and Gaussianity of primordial GWs (or gravitons) employing three-point correlators (equivalently, bispectra) of

* ohep2@cantab.ac.uk

† shiraishi_maresuke@rs.sus.ac.jp

the CMB temperature and polarization fields. This builds upon previous work including both a range of theoretical GW bispectrum predictions [10,23,33,35,38,39] and data analyses involving WMAP temperature (T) data [40] and *Planck* temperature and E -mode polarization (T and E) data [41,42], none of which have yet found any significant deviation from parity invariance and Gaussianity (see [43] for a review).¹

In this paper, we, for the first time,² include the *Planck* B -mode polarization (B) field in the data analysis (which dramatically enhances constraints on the parity-odd sector) and consider both general tests for non-Gaussianity and constraints on a specific inflationary template. This is facilitated by both new data (*Planck* PR4) with better control of large-scale polarization systematics and a suite of new CMB polyspectral estimators [56–58], which allow for quasioptimal estimation of binned parity-even and parity-odd T -, E -, and B -mode bispectra, while simultaneously accounting for observational effects such as mask-induced leakage between bins and polarizations. These have been previously used to assess parity violation in the scalar CMB sector for both temperature and polarization [51,52].

Parity-even and parity-odd information on graviton non-Gaussianity can be easily distinguished by working in harmonic space, where, for the CMB bispectrum, there is a bifurcation into even and odd $\ell_1 + \ell_2 + \ell_3$ [23,54,59]. Using the aforementioned recently developed general bispectrum estimator [56–58] (building on previous binned estimators [60,61] and numerical tricks [62,63]), we separately extract the parity-even and parity-odd information from CMB bispectrum data sets and perform model-independent tests of parity violation using the GW three-point function, essentially asking whether we detect any signal above the noise. Following this, we place constraints on the amplitude of a well-examined GW bispectrum template [39,41,43], parametrized by the non-Gaussianity parameter $f_{\text{NL}}^{\text{tens}}$ [see (10) for a concrete form]. This template can be defined for both the parity-even and parity-odd sectors and has a dominant signal for equilateral momentum configurations $k_1 \sim k_2 \sim k_3$. As such, it can well describe non-Gaussian chiral GWs sourced by gauge fields coupled to axions in various ways [9,10,32]. We place constraints on $f_{\text{NL}}^{\text{tens}}$ from both parity-even and -odd information alone as well as the combined (maximally chiral) limit; our results are consistent with zero apart from a $\simeq 2.5\sigma$ parity-even signal extracted from E -modes or their combination with B -modes. The most stringent value is obtained from the full T , E and B data set, with $f_{\text{NL}}^{\text{tens}} = 900 \pm 700$, which represents an improvement over

previous constraints ($f_{\text{NL}}^{\text{tens}} = 1300 \pm 1200$ [42]) by a factor greater than 2, mostly owing to the inclusion of B -modes.³

The remainder of this paper is structured as follows. In Sec. II we present our theoretical setup and motivations for testing parity and Gaussianity of the graviton sector. Section III discusses the *Planck* data and various aspects of its analysis, including bispectrum estimation and simulations. In Sec. IV we present our main result—constraints on the tensor amplitude parameter $f_{\text{NL}}^{\text{tens}}$ —before concluding in Sec. V. Throughout this work, we assume a *Planck* 2018 fiducial cosmology with the parameter set $\{\omega_b = 0.022383, \omega_{\text{cdm}} = 0.12011, h = 0.6732, \tau_{\text{reio}} = 0.0543, \log 10^{10} A_s = 3.0448, n_s = 0.96605\}$ for a single massive neutrino with $m_\nu = 0.06$ eV [1]. To avoid confirmation bias, the bispectrum estimation and analysis pipeline was finalized before the *Planck* bispectrum was computed.

II. THEORY

We begin by outlining the theoretical fundamentals of graviton non-Gaussianity, connecting them to CMB statistics, and discussing the specific CMB template whose amplitude will be constrained in Sec. IV C. Here and henceforth, we ignore correlations between the scalar and tensor sectors; the corresponding mixed non-Gaussianity represents an interesting extension of the ideas considered in this work.

A. Primordial tensors

In the synchronous gauge, we may write the metric perturbation induced by a graviton as $h_{ij} \equiv \delta g_{ij}^T / a^2$, where a is the scale factor and $\delta g_{\mu\nu}^T$ is the tensor part of the metric perturbation (at linear order). This can be decomposed into helicity- ± 2 states (equivalent to circular R/L polarizations) via

$$h_{ij}(\mathbf{x}) = \int_{\mathbf{k}} \sum_{\lambda=\pm 2} h_{\mathbf{k}}^{(\lambda)} e_{ij}^{(\lambda)}(\hat{\mathbf{k}}) e^{i\mathbf{k}\cdot\mathbf{x}} \quad (1)$$

for transverse-traceless helicity state $e_{ij}^{(\lambda)}(\hat{\mathbf{k}})$, obeying

$$e_{ij}^{(\lambda)}(-\hat{\mathbf{k}}) = e_{ij}^{(-\lambda)}(\hat{\mathbf{k}}) = e_{ij}^{(\lambda)*}(\hat{\mathbf{k}}), \quad e_{ij}^{(\lambda)}(\hat{\mathbf{k}}) e^{(\lambda'),ij}(\hat{\mathbf{k}}) = 2\delta_{\mathbf{k}}^{\lambda,-\lambda'} \quad (2)$$

[64], with $\int_{\mathbf{k}} \equiv (2\pi)^{-3} \int d^3\mathbf{k}$ throughout. The mode functions $h_{\mathbf{k}}^{(\lambda)}$ set the statistical properties of the graviton, with the three-point function given by

¹For the scalar sector, parity violation appears first as an imaginary component of the four-point functions (in the limit of negligible isotropy violation [44]). See [45–52] for searches for this signal.

²[53] also analyzed the B -mode polarization data, but only in the context of galactic foregrounds. Furthermore, [36,54,55] presented a detailed discussion of B -mode bispectra in the context of template analyses but did not apply their methodology to data.

³This $f_{\text{NL}}^{\text{tens}}$ is equivalent to f_{NL}^{P} in [40], $f_{\text{NL}}^{\text{tens}}$ in [41,42], and $f_{\text{NL}}^{\text{tens,eq}}$ in [43].

$$(2\pi)^3 \delta_{\text{D}}(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) B_h^{(\lambda_1 \lambda_2 \lambda_3)}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3) \equiv \left\langle \prod_{n=1}^3 h_{\mathbf{k}_n}^{(\lambda_n)} \right\rangle. \quad (3)$$

Under a conjugation, $h_{\mathbf{k}}^{(\lambda)} \rightarrow h_{\mathbf{k}}^{(\lambda)*} \equiv h_{-\mathbf{k}}^{(\lambda)}$, and thus

$$B_h^{(\lambda_1 \lambda_2 \lambda_3)*}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3) = B_h^{(\lambda_1 \lambda_2 \lambda_3)}(-\mathbf{k}_1, -\mathbf{k}_2, -\mathbf{k}_3). \quad (4)$$

A parity transform sends $\mathbf{x} \rightarrow \mathbb{P}[\mathbf{x}] \equiv -\mathbf{x}$; in this case, the helicity-state transforms as $\mathbb{P}[h_{\mathbf{k}}^{(\lambda)}] \rightarrow h_{-\mathbf{k}}^{(-\lambda)}$, and

$$\begin{aligned} \mathbb{P}[B_h^{(\lambda_1 \lambda_2 \lambda_3)}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3)] &= B_h^{(-\lambda_1 -\lambda_2 -\lambda_3)}(-\mathbf{k}_1, -\mathbf{k}_2, -\mathbf{k}_3) \\ &\equiv [B_h^{(\lambda_1 \lambda_2 \lambda_3)}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3)]^*. \end{aligned} \quad (5)$$

We can further decompose the bispectrum into parity-even/-odd components as $B_h^{(\lambda_1 \lambda_2 \lambda_3)}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3) = B_h^{(\lambda_1 \lambda_2 \lambda_3),+}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3) + B_h^{(\lambda_1 \lambda_2 \lambda_3),-}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3)$, which satisfy

$$\mathbb{P}[B_h^{(\lambda_1 \lambda_2 \lambda_3),\pm}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3)] = \pm B_h^{(\lambda_1 \lambda_2 \lambda_3),\pm}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3); \quad (6)$$

from (5), these have the symmetry $B_h^{(\lambda_1 \lambda_2 \lambda_3),\pm}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3) = \pm B_h^{(-\lambda_1 -\lambda_2 -\lambda_3),\pm}(-\mathbf{k}_1, -\mathbf{k}_2, -\mathbf{k}_3)$.

B. CMB fluctuations

Large-scale tensor modes source both CMB temperature and polarization fluctuations, which are here denoted by the harmonic coefficients $a_{\ell m}^X$, with $X \in \{T, E, B\}$. At linear order,

$$\begin{aligned} a_{\ell m}^X &= 4\pi i^\ell \int_{\mathbf{k}} \mathcal{T}_\ell^X(k) [h_{\mathbf{k}}^{(+2)} {}_{-2}Y_{\ell m}^*(\hat{\mathbf{k}}) + (-1)^x h_{\mathbf{k}}^{(-2)} {}_{+2}Y_{\ell m}^*(\hat{\mathbf{k}})] \\ &\equiv 4\pi i^\ell \int_{\mathbf{k}} \mathcal{T}_\ell^X(k) [h_{\mathbf{k}}^{(+2)} + (-1)^{x+\ell} h_{-\mathbf{k}}^{(-2)}] {}_{-2}Y_{\ell m}^*(\hat{\mathbf{k}}) \end{aligned} \quad (7)$$

[64,65], where \mathcal{T}_ℓ^X is the tensor-mode linear transfer function, ${}_s Y_{\ell m}$ is a spin-weighted spherical harmonic, and x encodes the parity of the field with 0 for T, E and 1 for B .

The angular bispectrum follows directly from (7) and (3),

$$\begin{aligned} \left\langle \prod_{n=1}^3 a_{\ell_n m_n}^{X_n} \right\rangle &= \prod_{n=1}^3 \left[4\pi i^{\ell_n} \int_{\mathbf{k}_n} \mathcal{T}_{\ell_n}^{X_n}(k_n) {}_{-2}Y_{\ell_n m_n}^*(\hat{\mathbf{k}}_n) \right] (2\pi)^3 \delta_{\text{D}}(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) \\ &\times \left\{ B_h^{(222)}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3) + (-1)^{\ell_{123} + x_{123}} B_h^{(-2-2-2)}(-\mathbf{k}_1, -\mathbf{k}_2, -\mathbf{k}_3) \right. \\ &+ (-1)^{x_1 + \ell_1} B_h^{(-222)}(-\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3) + (-1)^{x_{23} + \ell_{23}} B_h^{(2-2-2)}(\mathbf{k}_1, -\mathbf{k}_2, -\mathbf{k}_3) \\ &+ (-1)^{x_2 + \ell_2} B_h^{(2-22)}(\mathbf{k}_1, -\mathbf{k}_2, \mathbf{k}_3) + (-1)^{x_{13} + \ell_{13}} B_h^{(-22-2)}(-\mathbf{k}_1, \mathbf{k}_2, -\mathbf{k}_3) \\ &\left. + (-1)^{x_3 + \ell_3} B_h^{(22-2)}(\mathbf{k}_1, \mathbf{k}_2, -\mathbf{k}_3) + (-1)^{x_{12} + \ell_{12}} B_h^{(-2-22)}(-\mathbf{k}_1, -\mathbf{k}_2, \mathbf{k}_3) \right\} \end{aligned} \quad (8)$$

for $\ell_{ab} \equiv \ell_a + \ell_b$, etc. Splitting into parity-even and -odd graviton bispectra via (6), this simplifies considerably,

$$\begin{aligned} \left\langle \prod_{n=1}^3 a_{\ell_n m_n}^{X_n} \right\rangle &= \sum_{\pm} [1 \pm (-1)^{\ell_{123} + x_{123}}] \prod_{n=1}^3 \left[4\pi i^{\ell_n} \int_{\mathbf{k}_n} \mathcal{T}_{\ell_n}^{X_n}(k_n) {}_{-2}Y_{\ell_n m_n}^*(\hat{\mathbf{k}}_n) \right] (2\pi)^3 \delta_{\text{D}}(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) \\ &\times \left\{ B_h^{(222),\pm}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3) + (-1)^{x_1 + \ell_1} B_h^{(-222),\pm}(-\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3) \right. \\ &\left. + (-1)^{x_2 + \ell_2} B_h^{(2-22),\pm}(\mathbf{k}_1, -\mathbf{k}_2, \mathbf{k}_3) + (-1)^{x_3 + \ell_3} B_h^{(22-2),\pm}(\mathbf{k}_1, \mathbf{k}_2, -\mathbf{k}_3) \right\}. \end{aligned} \quad (9)$$

From the prefactor it is evident that parity-even (-odd) GW bispectra B_h^\pm source a nonvanishing CMB bispectrum signal only with $\ell_1 + \ell_2 + \ell_3 + x_1 + x_2 + x_3 = \text{even}$ (odd) [23,66]. In other words:

- (1) Parity-even GW bispectra B_h^+ source nonzero CMB bispectra with even $\ell_1 + \ell_2 + \ell_3$ if there are zero or two B -modes, and odd $\ell_1 + \ell_2 + \ell_3$ otherwise.
- (2) Parity-odd GW bispectra B_h^- source nonzero CMB bispectra with odd $\ell_1 + \ell_2 + \ell_3$ if there are

zero or two B -modes, and even $\ell_1 + \ell_2 + \ell_3$ otherwise.

This delineation is slightly more complex than in previous analyses [40–42] due to the presence of parity-odd B -modes. In Sec. III we will measure the parity-even and parity-odd signals separately on the basis of this rule and test the parity symmetry of the graviton bispectrum. Thus, a nonzero signal in a parity-odd correlator would give evidence of parity-violating non-Gaussianity in the graviton sector.

C. A well-motivated template

To perform model-specific CMB GW analyses, one needs to specify the form of the graviton three-point function $B_h^{(\lambda_1\lambda_2\lambda_3)}$. In this work, we consider a commonly used phenomenological template specified by

$$B_h^{(\lambda_1\lambda_2\lambda_3)}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3) = \frac{16\sqrt{2}}{27} f_{\text{NL}}^{\text{ttt}} \times B_\zeta^{\text{eq}}(k_1, k_2, k_3) e_{ij}^{(-\lambda_1)}(\hat{\mathbf{k}}_1) e_{jk}^{(-\lambda_2)}(\hat{\mathbf{k}}_2) \times e_{ki}^{(-\lambda_3)}(\hat{\mathbf{k}}_3) \delta_{\lambda_1+2}^{\text{K}} \delta_{\lambda_2+2}^{\text{K}} \delta_{\lambda_3+2}^{\text{K}} \quad (10)$$

[39–43], where B_ζ^{eq} is the standard equilateral-type scalar bispectrum, defined by

$$B_\zeta^{\text{eq}}(k_1, k_2, k_3) \equiv \frac{18}{5} (2\pi^2 A_s)^2 \left[- \left(\frac{1}{k_1^3 k_2^3} + 2 \text{ perms} \right) - \frac{2}{k_1^2 k_2^2 k_3^2} + \left(\frac{1}{k_1 k_2^2 k_3^3} + 5 \text{ perms} \right) \right], \quad (11)$$

assuming scale invariance. The bispectrum template of (10) contains only positive-helicity modes; as such, it is maximally chiral, containing both parity-even and parity-odd components with the same primordial amplitude, such that $B_h^{(\lambda_1\lambda_2\lambda_3),+}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3) = B_h^{(\lambda_1\lambda_2\lambda_3),-}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3) = (1/2) B_h^{(\lambda_1\lambda_2\lambda_3)}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3)$.⁴ A dominant signal arises in the equilateral configurations $k_1 \simeq k_2 \simeq k_3$. Notably, this is just one choice of template; one could also choose other forms such as local-type tensor non-Gaussianity (sourced, for example, by magnetic fields) or that induced by Weyl gravity [23,38,43].

The GW bispectra encoded by the above template can be produced in inflationary models that involve couplings between axions (χ) and gauge fields (A) through the Lagrangian

$$\mathcal{L}(A_\mu, \chi) \supset -\frac{1}{4} F_{\mu\nu} F^{\mu\nu} - \frac{\chi}{4f} F_{\mu\nu} \tilde{F}^{\mu\nu} \quad (12)$$

[7,8] [displayed for $U(1)$ gauge fields, though higher-order symmetries are also possible]. Here, $F_{\mu\nu} \equiv 2\partial_{[\mu} A_{\nu]}$ is the gauge field strength, $\tilde{F}_{\mu\nu}$ is its dual, and f is the Chern-Simons coupling parameter [5,7–11,32,67,68]. In such models, one of the two gauge field helicity states experiences a tachyonic instability and is explosively produced for Fourier modes with physical wavelengths comparable to the horizon. Through the gravitational interaction, the gauge field quadratically sources metric perturbations (with $h \sim A^2$); thus, supposing the Gaussianity of A_μ , the generated GW becomes chi-squared distributed and hence

⁴This can be split into distinct parity-even and parity-odd components through (6), which can be analyzed separately.

highly non-Gaussian. In this process, the chirality of the gauge field is transmitted to the GW. Without loss of generality, we identify the growing mode with $h^{(+2)}$, allowing us to drop any correlators involving $h^{(-2)}$.

Due to the gauge field production for limited Fourier modes, the chiral GW bispectrum has a dominant signal when the three bispectrum wave numbers are approximately equal, i.e., it peaks in equilateral configurations with $k_1 \simeq k_2 \simeq k_3$.⁵ Moreover, if the axion rolls slowly down a not-so-steep potential for a not-so-small period, the GW bispectrum can be shown to take the nearly scale-invariant shape of (10).⁶ In cases where the gauge field has $U(1)$ symmetry, the tensor nonlinearity parameter is evaluated as

$$f_{\text{NL}}^{\text{ttt}}|_{U(1)} \simeq 6.4 \times 10^{11} A_s^3 e^3 \frac{e^{6\pi\xi}}{\xi^9}, \quad (13)$$

where e is the inflaton slow-roll parameter and $\xi \equiv \dot{\chi}/(2fH)$, with $\dot{\chi}$ being the time derivative of the axion field and H the Hubble parameter [9,41]. In models involving an $SU(2)$ gauge field, the size of non-Gaussianity is determined by the energy density fraction of the gauge field Ω_A and the tensor-to-scalar ratio r as [32]

$$f_{\text{NL}}^{\text{ttt}}|_{SU(2)} \simeq 2.5 \frac{r^2}{\Omega_A}. \quad (14)$$

In Sec. IV C we estimate $f_{\text{NL}}^{\text{ttt}}$ from *Planck* T -, E -, and B -modes by inserting the above template into the bispectrum relation of (9). We give results for both the parity-even/odd signal (denoting the respective amplitudes as $f_{\text{NL}}^{\text{ttt},+/-}$) and the combined one (with amplitude $f_{\text{NL}}^{\text{ttt}} \equiv f_{\text{NL}}^{\text{ttt},+} + f_{\text{NL}}^{\text{ttt},-}$). For this purpose, we may restrict the range of CMB multipoles to $\ell \in [2, 500)$ as we expect no gain in the signal-to-noise ratio when including $\ell \gtrsim 500$ due to the rapid damping of CMB tensor modes [39,43].

III. DATA SET AND ANALYSIS PIPELINE

A. Observational data

In this work, we obtain constraints on GW bispectra using the latest temperature and polarization maps provided by *Planck*. Our data set is analogous to that used in [52]; we use the PR4 data set, consisting of component-separated temperature and Q/U -mode polarization maps processed using the NPIPE pipeline in combination with the SEVEM

⁵If one instead considers a Lagrangian of the form $\mathcal{L} \supset f(\chi)(-\frac{1}{4}F^2 + \frac{1}{4}F\tilde{F})$ and appropriately fine-tunes the time evolution of the coupling function $f(\chi)$, one can produce a gauge field that does not decay on superhorizon scales. In this case, a sizable bispectrum signal can then arise in the squeezed limit: $k_1 \sim k_2 \gg k_3$, $k_2 \sim k_3 \gg k_1$ and $k_3 \sim k_1 \gg k_2$ [69].

⁶See [10,11] for cases inducing signals with strong scale dependence.

component separation algorithm [70–72].⁷ These are an updated version of the *Planck* 2018 data set (PR3), and have a number of improvements, in particular with regards to the treatment of large-scale polarization data and reduction of systematics. We additionally utilize 600 FFP10 simulations (all processed with NPIPE); 100 are used to construct the quasioptimal bispectrum estimators and the remainder are used for validation and to estimate the variance of the observed statistics [70].

Following [42], high-emission regions of the Galactic plane are removed via the *Planck* common component-separation mask (smoothed at ten arcminute scales), which preserves $\simeq 80\%$ of the sky [73]. In the analyses below, we additionally apply the *Planck* beam, which includes the PR4 polarization transfer function on large scales ($\ell < 40$), extracted from the FFP10 simulations [70]. Since the GW signatures we consider in this study are predominantly confined to large angular scales, we restrict our attention to $\ell_{\max} = 500$, choosing a conservative HEALPix pixelation strategy with $N_{\text{side}} = 512$ [74] for all maps.

Before computing bispectra (Sec. III B), all maps are filtered by a linear operator, \mathbf{S}^{-1} [cf. [62,63]]. As discussed in [56,57], the choice of \mathbf{S}^{-1} sets the precision of the output CMB correlators, with optimal variances achieved if \mathbf{S}^{-1} is equal to the true inverse pixel covariance of the data set. In practice, such a choice is difficult to implement, and thus here we adopt a simpler scheme following [52]. Introducing a mask $W(\hat{\mathbf{n}})$ and a linear inpainting operator Π , the action of \mathbf{S}^{-1} on some map $a^X(\hat{\mathbf{n}})$ (for $X \in \{T, E, B\}$) can be written as

$$[\mathbf{S}^{-1}a]_{\ell m}^X = \sum_Y S_{\ell}^{-1,XY} \int d\hat{\mathbf{n}} Y_{\ell m}^*(\hat{\mathbf{n}}) W(\hat{\mathbf{n}}) \Pi[a^Y](\hat{\mathbf{n}}), \quad (15)$$

i.e., we inpaint small holes in the map (cf. [75]), mask the Galactic plane, transform to harmonic space, and then normalize by a rotationally invariant filter. Here, we fix $S_{\ell}^{XY} = C_{\ell}^{XY} + \delta_{\mathbf{K}}^{XY} N_{\ell}^X$ with fiducial CMB power spectrum C_{ℓ}^{XY} and diagonal noise N_{ℓ}^X , as measured from the high-resolution *Planck* half-mission maps. We stress that our choice of \mathbf{S}^{-1} filter cannot induce bias; it is fully accounted for in the bispectrum estimator discussed below and any simplification just leads to a reduction in sensitivity.

B. Bispectrum estimation

In the remainder of this work, we consider only *reduced bispectra* $b_{\ell_1 \ell_2 \ell_3}^{XYZ}$, which are defined from CMB maps $a_{\ell m}^X$ via

$$\begin{aligned} & \langle a_{\ell_1 m_1}^X a_{\ell_2 m_2}^Y a_{\ell_3 m_3}^Z \rangle_c \\ &= \frac{1}{3} \sqrt{\frac{(2\ell_1 + 1)(2\ell_2 + 1)(2\ell_3 + 1)}{4\pi}} \\ & \times \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ m_1 & m_2 & m_3 \end{pmatrix} \left[\begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ -1 & -1 & 2 \end{pmatrix} + 2 \text{ perms} \right] b_{\ell_1 \ell_2 \ell_3}^{XYZ} \\ & \equiv w_{m_1 m_2 m_3}^{\ell_1 \ell_2 \ell_3} b_{\ell_1 \ell_2 \ell_3}^{XYZ}, \end{aligned} \quad (16)$$

where the weights $w_{m_1 m_2 m_3}^{\ell_1 \ell_2 \ell_3}$ are a generalization of the Gaunt integral, allowing for odd $\ell_1 + \ell_2 + \ell_3$ [54,57]. This restricts to rotationally invariant bispectra (due to the first Wigner $3j$ symbol), reducing the dimensionality from six to three angular variables. Conjugation symmetries of $a_{\ell m}^X$ imply that bispectra with even (odd) $\ell_1 + \ell_2 + \ell_3$ are real (imaginary).

To estimate bispectra, we use the PolyBin code⁸ described in [56,57] (which builds upon template-based techniques from e.g., [62,63], the binned estimators of [60,61], and the three-dimensional estimators of [76,77]). Rather than computing the full bispectrum [which has dimensionality $\mathcal{O}(\ell_{\max}^3)$], this directly estimates the bispectrum in some set of angular bins, $\mathbf{b} \equiv \{b_1, b_2, b_3\}$, hereafter denoted $b_{\chi}^{XYZ}(\mathbf{b})$, where $\chi = \pm 1$ indicates the parity of the correlator. In more detail:

- (1) *Parity-even*, $\chi = +1$: this corresponds to *even* $\ell_1 + \ell_2 + \ell_3$ and *real* bispectra for correlators with zero or two *B*-modes, and *odd* $\ell_1 + \ell_2 + \ell_3$ and *imaginary* bispectra otherwise.
- (2) *Parity-odd*, $\chi = -1$: this corresponds to *odd* $\ell_1 + \ell_2 + \ell_3$ and *imaginary* bispectra for correlators with zero or two *B*-modes, and *even* $\ell_1 + \ell_2 + \ell_3$ and *real* bispectra otherwise.

Viz., the discussion in Sec. II C, parity-even (parity-odd) inflationary templates such as $f_{\text{NL}}^{\text{tt},+}$ ($f_{\text{NL}}^{\text{tt},-}$) will correspond to $\chi = +1$ ($\chi = -1$).

Schematically, PolyBin estimates bispectra using two components: 1) a numerator, $[\mathcal{F}\hat{b}]$, involving three copies of the data and 2) a data-independent normalization matrix, \mathcal{F} (often called the Fisher matrix).⁹ As discussed in [56,57], this is an unbiased estimator for any choice of weighting scheme, implying that the output spectra can be compared to theory without the need to account for window-function convolution or leakage between different bins, parities, and fields. Furthermore, in the ideal limit, where \mathbf{S}^{-1} is equal to the inverse pixel covariance and $a_{\ell m}^X$ is Gaussian, the estimator is optimal and has covariance \mathcal{F}^{-1} (e.g., [78]).

⁸Available at github.com/OliverPhilcox/PolyBin [58].

⁹When forming the estimator, one has freedom in where to include the mask, $W(\hat{\mathbf{n}})$. Following [51,52], here we include it in the \mathbf{S}^{-1} weights rather than directly in the estimator (as in [56,57]) since this (empirically) reduces bin-to-bin leakage without inducing bias.

⁷Publicly available at portal.nersc.gov/project/cmb/planck2020.

Given beam B_ℓ^X and a triplet of fields $u \equiv \{XYZ\}$ containing n_B B -modes, the estimator numerator takes the schematic form

$$\begin{aligned}
 [\mathcal{F}\hat{b}]_\chi^u(\mathbf{b}) &\sim \sum_{\ell_i \in b_i} \sum_{m_i} [1 + \chi(-1)^{\ell_1 + \ell_2 + \ell_3 + n_B}] \\
 &\times B_{\ell_1}^{u_1} B_{\ell_2}^{u_2} B_{\ell_3}^{u_3} W_{m_1 m_2 m_3}^{\ell_1 \ell_2 \ell_3} \\
 &\times [a_{\ell_1 m_1}^{u_1} a_{\ell_1 m_1}^{u_1} a_{\ell_1 m_1}^{u_1} - 3 \langle a_{\ell_1 m_1}^{u_1} a_{\ell_2 m_2}^{u_2} \rangle a_{\ell_3 m_3}^{u_3}], \quad (17)
 \end{aligned}$$

summing over all ℓ modes in the bin of interest. This involves both a cubic and a linear term; the latter does not change the mean of the estimator but can considerably reduce the variance on large scales. The normalization matrix has a more complex form, which is set by demanding that the estimator be unbiased [such that $\mathbb{E}[\hat{b}_\chi^u(\mathbf{b})] = b_\chi^u(\mathbf{b})$ for expectation operator \mathbb{E}]. In practice, PolyBin does not compute the $\mathcal{O}(\ell_{\max}^6)$ summations in the above expressions directly: efficient use of spherical-harmonic transforms and Monte Carlo summation allows both the numerator and denominator to be efficiently computed with $\mathcal{O}(\ell_{\max}^3)$ complexity. All computations are done in PYTHON, with spherical harmonic transforms performed using Libsharp [79].

In this work, we compute bispectra for each of the ten nontrivial combinations of the T, E, B fields:

$$\{TTT, TTE, TTB, TEE, TEB, TBB, \\
 EEE, EEB, EBB, BBB\}, \quad (18)$$

with $\chi \in \{\pm 1\}$. Based on the discussion in Sec. II C, we restrict to $\ell_{\max} = 500$; noting that the computation time of the estimator denominator is cubic in the number of ℓ bins, N_ℓ (but linear in the total number of bins, N_{bins}), we use $N_\ell = 13\ell$ bins roughly equally spaced in $\ell^{2/3}$ (ensuring roughly equal signal-to-noise per bin), with bin edges

$\{2, 3, 4, 22, 49, 82, 121, 164, 211, 262, 317, 375, 436, 500\}$. This gives a total of $N_{\text{bins}} = 5972$ bispectrum bins (dropping any empty bins and removing degeneracies by restricting to $b_1 \leq b_2 \leq b_3$ for TTT , etc.).¹⁰ To form the linear term of the bispectrum estimator (17), we require the expectation $\langle a_{\ell_1 m_1}^{u_1} a_{\ell_2 m_2}^{u_2} \rangle$; rather than computing this analytically (which is expensive, since the mask breaks rotational invariance), we estimate the entire linear term as a Monte Carlo average of the data and $N_{\text{sim}} = 100$ FFP10 simulations [requiring $\mathcal{O}(N_{\text{sim}} N_\ell)$ computational costs, if harmonic transforms are rate limiting], as discussed in [57]. Furthermore, the normalization matrix is also computed as a Monte Carlo average; due to the efficient numerical schemes introduced in [62], we are able to estimate the matrix to percent-level precision using only $N_{\text{mc}} = 10$ realizations [with $\mathcal{O}(N_{\text{mc}} N_{\text{bin}})$ computational cost].

Using the above hyperparameters, computing the bispectrum numerator of a single data set requires 30 CPU-hours and around 8000 harmonic transforms (with 25 GB memory). This is dominated by the linear term; if this is removed (which we will later show yields only a minor inflation in our constraints), the computation time decreases to only 5 CPU-minutes with 80 transforms. The normalization matrix (which is independent of the data) requires 30 CPU-hours and 60 000 harmonic transforms per realization on a (partially utilized) high-performance node with 3 TB memory. In total, the full computation of the bispectra of 500 FFP10 simulations and the *Planck* data requires around 20 000 CPU-hours.

C. Binning the theoretical templates

To compare the data with bispectrum templates [such as (10)], one must recast the theory into the same form as the data. As discussed in [57], we can relate unbinned and binned correlators via the expectation of the idealized bispectrum estimator; this yields

$$\begin{aligned}
 [\mathcal{F}b]_\chi^{\text{th},u}(\mathbf{b}) &\propto \frac{1}{\Delta^u(\mathbf{b})} \sum_{\ell_1 \in b_1} \sum_{\ell_2 \in b_2} \sum_{\ell_3 \in b_3} \sqrt{\frac{(2\ell_1 + 1)(2\ell_2 + 1)(2\ell_3 + 1)}{4\pi}} \frac{1}{3} \left[\begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ -1 & -1 & 2 \end{pmatrix} + 2 \text{ perm} \right] \\
 &\times \sum_{u'} S_{\ell_1}^{-1, u_1 u_1'} S_{\ell_2}^{-1, u_2 u_2'} S_{\ell_3}^{-1, u_3 u_3'} B_{\ell_1 \ell_2 \ell_3}^{u_1' u_2' u_3'} \Big|_{\text{th}}, \\
 \mathcal{F}_{\chi\chi'}^{\text{th},u'u}(\mathbf{b}', \mathbf{b}) &\propto \frac{1}{\Delta^u(\mathbf{b}) \Delta^{u'}(\mathbf{b}')} \sum_{\ell_1 \in b_1} \sum_{\ell_2 \in b_2} \sum_{\ell_3 \in b_3} \frac{(2\ell_1 + 1)(2\ell_2 + 1)(2\ell_3 + 1)}{4\pi} \frac{1}{9} \left[\begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ -1 & -1 & 2 \end{pmatrix} + 2 \text{ perms} \right]^2 \\
 &\times [S_{\ell_1}^{-1, u_1 u_1'} S_{\ell_2}^{-1, u_2 u_2'} S_{\ell_3}^{-1, u_3 u_3'} \delta_{b_1' b_1}^K \delta_{b_2' b_2}^K \delta_{b_3' b_3}^K + 5 \text{ perms}], \quad (19)
 \end{aligned}$$

where we use weighting matrices S_ℓ^{XY} and assume that the theoretical bispectrum was specified by the isotropic form [which can be extracted from (9), noting that the Dirac delta can be recast as a $3j$ symbol, after angular integration]

¹⁰In the analysis of Sec. IV, we drop five of these bins, since their variance is both extremely large and poorly measured.

$$\langle a_{\ell_1 m_1}^X a_{\ell_2 m_2}^Y a_{\ell_3 m_3}^Z \rangle|_{\text{th}} = \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ m_1 & m_2 & m_3 \end{pmatrix} B_{\ell_1 \ell_2 \ell_3}^{XYZ} \Big|_{\text{th}}. \quad (20)$$

Here, the u' summations in $[\mathcal{F}b]^{\text{th}}$ extend over all 27 triplets of fields [not just those in (18)] and $\Delta^u(\mathbf{b})$ gives the degeneracy of the bin: 6 if $u_1 = u_2 = u_3$ and $b_1 = b_2 = b_3$, 2 if $u_1 = u_2$ and $b_1 = b_2$, 2 if $u_2 = u_3$ and $b_2 = b_3$, and 1 otherwise. Finally, we implicitly restrict to $(-1)^{\ell_1 + \ell_2 + \ell_3 + n_B} = \chi$ in each case and note that terms with $\chi \neq \chi'$ are independent, i.e., bispectra sourced by parity-even and parity-odd physics are independent. Here, the complexity of the normalization appears due to the nonzero TE cross spectra; if this vanishes, \mathcal{F}^{th} is diagonal, and the action of the above is just to appropriately average the theory across each ℓ bin. Note that one could alternatively drop the S^{-1} terms in both the estimator and the binned theory, computing an *unweighted* average over all modes in the bin; from a Fisher forecast, this was found to increase the error bar on T -mode parameter constraints by up to 45% (or 5% for B -modes), showing the utility of our weighting scheme.

D. Covariances and likelihoods

Finally, we consider how to constrain theoretical models (and non-Gaussianity amplitudes) from the data. As discussed in [52], it is necessary to make a theoretical ansatz for the covariance of the data set, $\hat{b}_\chi^u(\mathbf{b})$, since the full covariance has $\simeq 6000^2$ elements and cannot be robustly estimated from the 500 simulations. As in previous works, we assume that the correlation structure of $\text{Cov}(\hat{b}, \hat{b}')$ is well described by the inverse normalization matrix, \mathcal{F}^{-1} ; theoretically, this is appropriate since \mathcal{F}^{-1} is the covariance if the estimator is optimal and, empirically, it is found to be an excellent approximation, as demonstrated in Sec. III. In this approach, we first (reversibly) project the data and theory onto the Cholesky factorization of \mathcal{F} (e.g., [78]), defining

$$\beta_\chi^u(\mathbf{b}) \equiv [\mathcal{F}^{\text{T}/2} b]_\chi^u(\mathbf{b}) \quad (21)$$

and assume the following (diagonal) likelihood:

$$\mathcal{L}(\hat{b}|b^{\text{th}}) \propto \exp\left(-\frac{1}{2}\chi^2(\hat{b}|b^{\text{th}})\right), \quad (22)$$

$$\chi^2(\hat{b}|b^{\text{th}}) = \sum_{u,\chi,\mathbf{b}} \frac{[\hat{\beta}_\chi^u(\mathbf{b}) - \beta_\chi^{\text{th},u}(\mathbf{b})]^2}{\text{var}[\beta_\chi^u(\mathbf{b})]},$$

where b^{th} is some theoretical model and the variance is estimated from the FFP10 realizations. If one works with only a subset of the data, it is important to take this subset before the rotation is applied (i.e., to excise bins from \mathcal{F} and \hat{b} rather than from $\hat{\beta}$); otherwise, the results may incur bias. In the limit of an ideal estimator applied to Gaussian

data, $\text{var}[\beta_\chi^u(\mathbf{b})] = 1$; thus, deviations from unity will encode non-Gaussianity from residual foregrounds and suboptimality in the treatment of the mask.

The above likelihood allows for two tests of the data. First, we can perform blind tests to assess whether there is any evidence for a nonzero bispectrum of a given parity (cf. [48,51,52]). In this case, one sets $b^{\text{th}} = 0$ and compares the value of $\chi^2(\hat{b}|0)$ in the data (dropping bins of the wrong parity before projection) to the empirical distribution from simulations (or the theoretical expectation: a χ^2 distribution with N_{bins} degrees of freedom). Second, we can constrain individual theoretical models, such as the gauge-field template given in Sec. II C. Writing $b_\chi^{\text{th},u} = f_{\text{NL}}^{\text{th},u} b_\chi^{\text{template},u}$ for template b^{template} and amplitude $f_{\text{NL}}^{\text{th},u}$, constraints are obtained by sampling from the following posterior:

$$\mathcal{P}(f_{\text{NL}}^{\text{th},u}|\hat{b}) \propto \exp\left(-\frac{1}{2}\chi^2(\hat{b}|f_{\text{NL}}^{\text{th},u} b^{\text{template}})\right). \quad (23)$$

Often, one restricts to models of a given parity ± 1 , which can be achieved by replacing $f_{\text{NL}}^{\text{th},u} \rightarrow f_{\text{NL}}^{\text{th},\pm 1}$ and nulling bins in b^{template} with $\chi = \mp 1$.¹¹ We note that constraints obtained from a full theory model will generically be stronger than those from a blind test, since it allows for optimal model projection.

IV. RESULTS

We now present the main results of this paper: estimates of the *Planck* T -, E -, and B -mode bispectra and constraints on the tensor-mode non-Gaussianity. We first examine the data itself in Sec. IV A, before showing the model-independent blind-test constraints in Sec. IV B and the constraints on $f_{\text{NL}}^{\text{th}}$ in Sec. IV C.

A. Data

In Fig. 1, we display the measured *Planck* and FFP10 bispectra, projected as in (21). Since this data set contains almost 6000 bins, visual assessment is difficult; however, the *Planck* results do not appear to deviate wildly from the simulations, which themselves have averages highly consistent with zero (indicating that our pipeline does not contain significant additive bias). As mentioned above, if the data is Gaussian and the estimator is optimal, we would expect the variance of β to be unity; here, we find some departure from this, which depends strongly on the triangle configuration (leading to the spikes in the figure). In general, we find slightly enhanced variances for triangles containing polarized low- ℓ modes; this indicates that our polarization estimators are somewhat suboptimal, for example, due to residual non-Gaussian foregrounds

¹¹Since the rotation matrix can mix parities through the mask, this is not equivalent to restricting the summation of (22) to the relevant value of χ .




FIG. 1. Large-scale binned bispectrum measurements from *Planck* PR4 (blue points) compared to 500 FFP10 simulations (red bands) with $2 \leq \ell < 500$, computed with the PolyBin code. We show the data (top) and error (bottom) for 20 combinations of bins and parities, separated by dashed lines and indicated in green. Spectra marked + are parity symmetric, while those marked – are parity antisymmetric (with + indicating even $\ell_1 + \ell_2 + \ell_3 + n_B$ and – indicating the opposite, for n_B B -modes in the statistic). In each subpanel, all bispectrum bins are stacked, with the smallest scales on the left-hand side, and we take the imaginary part of any spectra with odd $\ell_1 + \ell_2 + \ell_3$. We plot the rotated bispectrum defined in (21); if the estimator is optimal and the data is Gaussian, this should be a unit normal variable. Here, we find slight enhancements in the variance (bottom panel), particularly for $\ell = 2$ modes and polarization, indicating the effects of large-scale non-Gaussianity and foreground contamination. This slightly enhances the variances of the parameter constraints, but does not lead to bias. Finally, the pink lines in the bottom panel show the effect of removing the linear term in the bispectrum estimator; we see a slight enhancement of the variance at low ℓ .

(as discussed below). Furthermore, we find that the large-scale variances are increased further if one drops the linear term in the bispectrum estimators [pink, cf. (17)]; this is as expected, and we consider its effects on the $f_{\text{NL}}^{\text{eff}}$ constraints in Sec. IV C. We stress that $\sigma(\beta) > 1$ does not imply that our estimator is biased, just that the variances will be slightly larger than would be obtained from an ideal analysis.

To more thoroughly assess the covariance of our data set, we turn to the correlation matrix of the (unprojected) bispectrum, defined as

$$\text{Corr}(\hat{b}_i, \hat{b}_j) = \frac{\text{Cov}(\hat{b}_i, \hat{b}_j)}{\sqrt{\text{Var}(\hat{b}_i)\text{Var}(\hat{b}_j)}}, \quad (24)$$

where i, j jointly indexes the angular bin, field, and parity. This is shown in the left panel of Fig. 2, for both the theory (inverting the PolyBin normalization matrix \mathcal{F}) and the FFP10 numerical covariance. In both cases, we find a roughly diagonal structure, but with nontrivial correlations at the $\simeq 2\%$ level between different fields, resulting from the

nonzero TE power spectrum. If one zooms in, smaller correlations (at the $\sim 0.5\%$ level) can be observed within the configuration of a given field (e.g., b^{TTT}); these arise from mask-induced leakage across ℓ bins. Although these effects are small (due to the broad bins used herein), they could lead to biases in our bispectrum estimates if uncorrected; in our case, the PolyBin normalization takes care of such complexities, returning an unbiased estimate of the binned spectrum.

In Sec. III D, we claimed that the inverse normalization matrix, \mathcal{F}^{-1} , could be used to decorrelate the bins in the bispectrum data vector, such that the projected statistic, β , would have diagonal covariance. From the left panel of Fig. 2, this assumption appears excellent, with the theoretical correlation structure appearing to be highly consistent with that of the mocks. This is quantified in the right panel of Fig. 2, where we instead plot the correlation structure of $\hat{\beta}$. As required, this is highly diagonal, with deviations fully consistent with noise,¹² implying that our

¹²The mean pixel value is -5.3×10^{-6} with a standard deviation of 0.045, against an expected value of 0.000 ± 0.045 .

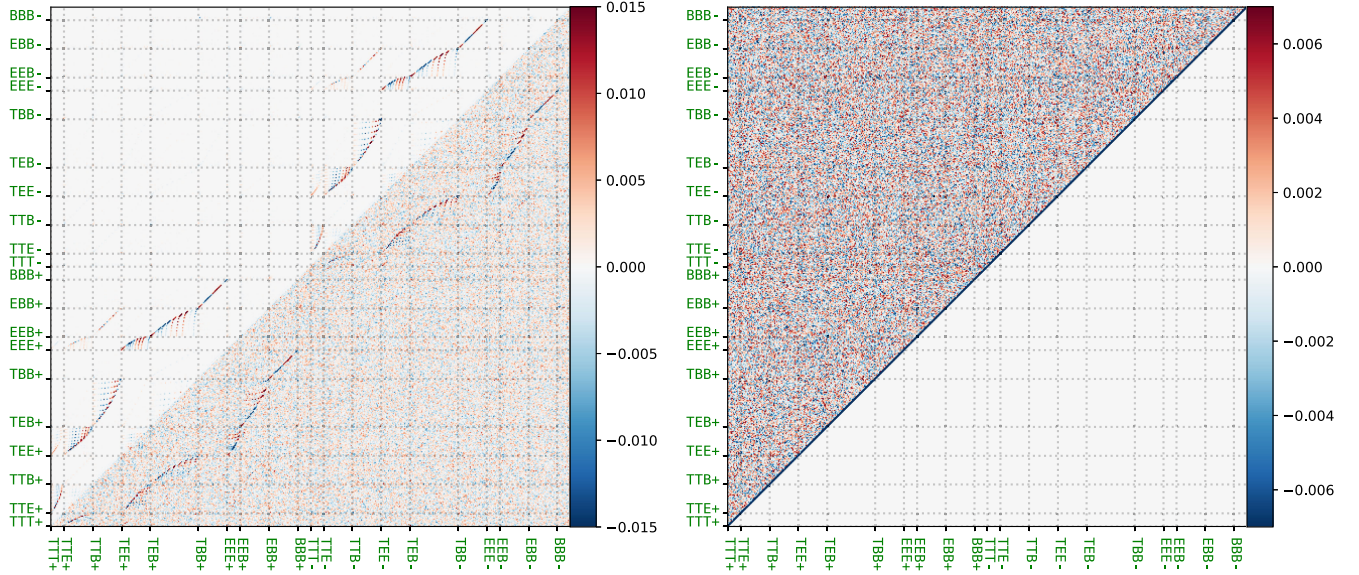


FIG. 2. Left: diagonal-subtracted correlation matrix of the binned bispectrum, $b_\chi^u(\mathbf{b})$, obtained from theory (inverting \mathcal{F} , top left), and from 500 FFP10 simulations (bottom right), organizing components as in Fig. 1. (Small) off-diagonal contributions arise from mask-induced leakage between different bins and parities and intrinsic correlations between T - and E -modes. Right: correlation matrix for the projected $\beta_\chi^u(\mathbf{b})$ statistic (21), with the (symmetric) lower half removed. We find excellent agreement between the theory and mock $b_\chi^u(\mathbf{b})$ covariance, indicating that the normalization matrix (which includes the mask and weighting scheme) captures the structure of the covariance well; this is evidenced by the covariance of τ , which appears entirely diagonal.




FIG. 3. Ratio of true to ideal inverse bispectrum variances (denoted $\Psi \equiv \text{diag}[C^{-1}]$, assuming $f_{\text{sky}} = 0.78$). The points show the ratios obtained from the FFP10 inverse variances, while the lines show those predicted by the inverse normalization matrix, \mathcal{F}^{-1} . Red and green points indicate triangles containing modes with $\ell = 2$ or $\ell = 3$, respectively. Nonuniformities in the mask strongly reduce the precision of squeezed bispectrum measurements, as shown by the theory curves (which include the window); such triangles, and B -mode spectra, are further significantly impacted by non-Gaussian foreground residuals present in the FFP10 covariance.

assumption of a diagonal likelihood (22) is robust, even in the presence of residual foregrounds found within FFP10.

Finally, we turn to the (unprojected) bispectrum variances. In Fig. 5 we plot the variance obtained from theory and simulations, finding excellent agreement between theory and simulations across many orders of magnitudes (as discussed above). Furthermore, we can assess the behavior of $\text{var}(\hat{b})$ for various bispectrum parameters: the behavior is broadly similar between the parity-even and -odd configurations, particularly at large ℓ , though a few bins have (almost) vanishingly small covariance, since very few modes can populate them while obeying the restrictions on the sign of $\ell_1 + \ell_2 + \ell_3$. Different fields also have different variance structures: this occurs due to the very different signal-plus-noise power spectra, with, for example, the B -modes being dominated by almost

scale-invariant noise. To assess the impact of masks and non-Gaussianity on this covariance, we show the ratio of the inverse covariance diagonal to that predicted from idealized theory (from the diagonal of the idealized \mathcal{F} given in (19) in Fig. 3. The PolyBin estimate for the covariance encodes the distortions induced by the mask, but assumes Gaussianity; here, we find that such distortions have a scale-dependent effect on the inverse variances, which lead to considerable degradation in the constraints on triangles containing small ℓ (at a similar level for each correlator). The FFP10 results also include non-Gaussian contributions from noise and foregrounds; these are seen to significantly inflate the covariance of B -modes and large-scale triangles. Such effects are difficult to remove, but will be lessened in the future by higher experimental sensitivity and better foreground deprojection schemes.

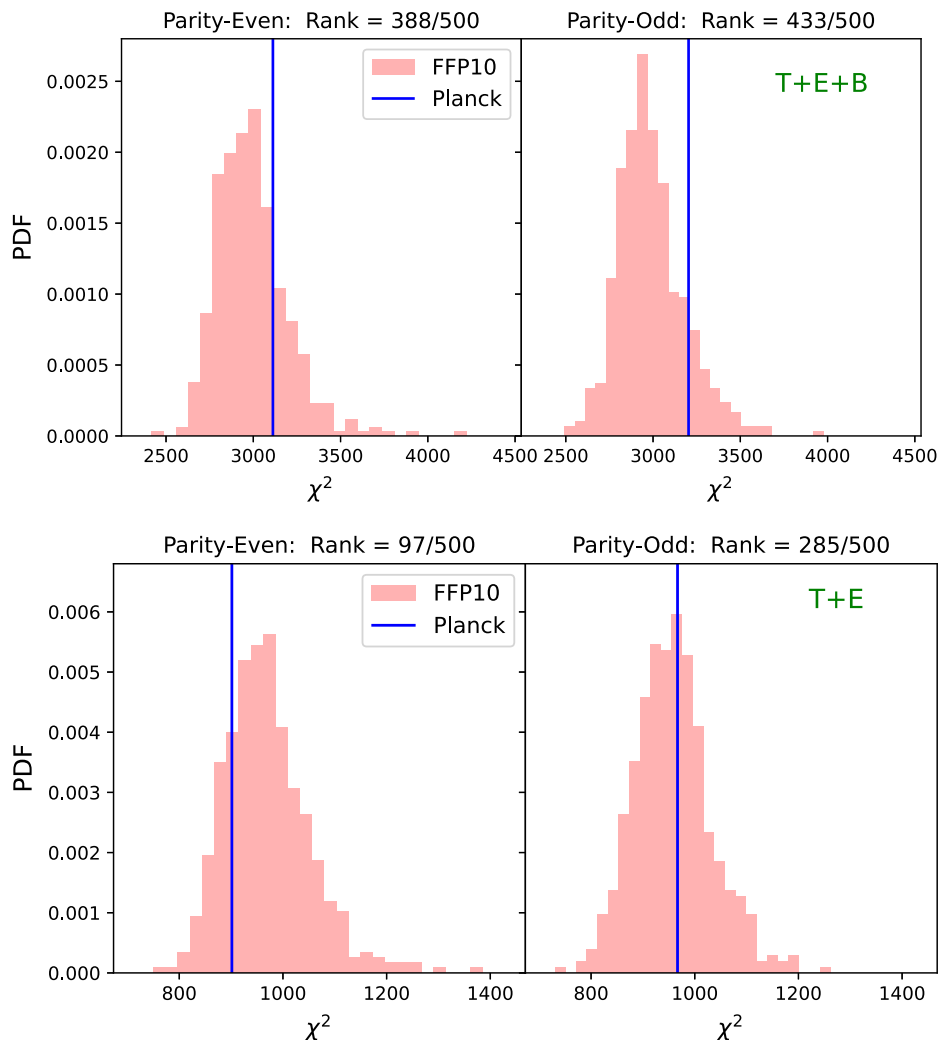


FIG. 4. Blind tests for bispectra from the *Planck* PR4 data set. The left and right panels show the constraints on parity-even and parity-odd bispectra, respectively, with the bottom panels displaying the results without B -modes. In all cases, we compute the χ^2 statistic (22), assuming zero signal, and assign a rank to the *Planck* data set (blue) based on the distribution of values from 500 FFP10 simulations (red), i.e., we perform a rank test. The *Planck* data is consistent with noise within 1.1σ ; we find no evidence for non-Gaussianity in this model-independent test.

B. Blind tests

Armed with the *Planck* data set and the FFP10 simulations, we may ask a more generic question: does our data set contain “any” evidence for bispectrum non-Gaussianity? To answer this, we perform blind tests similar to those used to parity violation in the CMB and large-scale structure four-point functions [48,51,52]. This follows the methodology summarized in Sec. III D; in brief, we compute a χ^2 parameter from the *Planck* data, relative to null assumptions, and compare its value to the empirical distribution extracted from simulations. For this purpose, we consider the parity-even and parity-odd sectors separately, and perform the test for a variety of data cuts, primarily focusing on $T + E + B$ and $T + E$ correlators (with the latter used in [42]).

The corresponding results are shown in Fig. 4 and Table I. When B -modes are included in the analysis, the empirical χ^2 distribution clearly departs from the (ideal) χ^2 distribution; for this reason, we do not compare the *Planck* data to a theoretical noise distribution. This likely arises from non-Gaussian residual foregrounds at low ℓ and is significantly reduced by removing B -modes or increasing ℓ_{\max} to 4. For the parity-even sector, the *Planck* data is consistent with the simulations at 0.8σ ; the result is similar $[(-)0.9\sigma]$ if B -modes are removed. This is an indirect constraint on both scalar and tensor non-Gaussianity, both of which can source parity-conserving physics. For parity-odd bispectra, we find consistency at 1.1σ (0.2σ) with (without) B -modes. This scenario is not sensitive to scalar physics; as argued in [80] (and many other works), parity-odd bispectra can be sourced only by vector and tensor-mode physics. These conclusions are stable across an array of analysis choices (though the precise significances vary due to differing noise realizations): we detect no signal in

TABLE I. Model-independent constraints on tensor non-Gaussianity and parity violation using *Planck* PR4 bispectrum data with various analysis choices. In each case, we compare the *Planck* χ^2 value [specified in (22)] to the empirical distribution obtained from 500 FFP10 simulations, performing a rank test. Results are given in equivalent Gaussian σ , noting that only positive detections are physical. The fiducial analysis assumes $\ell_{\min} = 2$ and $\ell_{\max} = 500$. All detection significances are below 2.3σ , indicating no compelling evidence for new physics.

Analysis	Parity-even	Parity-odd
Fiducial ($T + E + B$)	0.8σ	1.1σ
T only	-0.4σ	-1.2σ
E only	-1.2σ	0.4σ
B only	1.7σ	1.5σ
$T + E$	-0.9σ	0.2σ
$T + B$	0.4σ	1.0σ
$E + B$	1.2σ	1.3σ
$\ell_{\min} = 4$	1.8σ	2.3σ
$\ell_{\max} = 375$	0.8σ	1.3σ

any combination of T -, E -, and B -modes and the results are broadly insensitive to scale cuts, with a maximum detection significance of 2.3σ . As such, *we find no compelling evidence for tensor non-Gaussianity or parity violation.*

Whilst we do not detect any signal in this study, we should make clear that this is *not* an optimal test for specific models of non-Gaussianity; a better constraint can always be obtained in a model-specific analysis, as we perform below for the template given in (10). However, since the current work is the first to consider generic B -mode bispectra from observational data, it is pertinent to perform a general study in addition to model-specific constraints.

C. Model constraints

1. Theoretical spectra

Finally, we obtain constraints on the graviton bispectrum template specified in (10). Before constraining the amplitude parameter $f_{\text{NL}}^{\text{tll}}$, we briefly discuss the form of the binned spectra. As shown in Fig. 5, both parity-even and parity-odd components of the theory source real and imaginary parts; this is due to the presence of B -modes and follows the logic described in Sec. III B. We find strong scale dependence of the bispectrum signal, with the largest amplitudes observed for the lowest ℓ modes. There is some dependence on field, with B -modes having a lower amplitude; however, B -mode spectra do not suffer from cosmic variance, and thus their inclusion is expected to considerably sharpen constraints on $f_{\text{NL}}^{\text{tll},\pm}$. Finally, we note similar forms for the parity-odd and parity-even physics; this is as expected, since the total number of modes with odd and even $\ell_1 + \ell_2 + \ell_3$ are similar. In contrast the variance of modes with even and odd $\ell_1 + \ell_2 + \ell_3$ can differ considerably (e.g., $EEE+$ has much lower variance than $EEE-$, with the opposite behavior for $BBB+$ and $BBB-$); this is due to the Wigner $3j$ factors (which cancel in the theory model, but not in the bispectrum variance), which asymptotically vanish for $\ell_1 + \ell_2 + \ell_3$, i.e., any signal is eventually washed out by projection effects.

2. Main constraints

In Fig. 6 and Table II, we present constraints on the tensor non-Gaussianity amplitudes $f_{\text{NL}}^{\text{tll},\pm}$ from *Planck* T , E and (optionally) B -mode data. In our fiducial analysis, including all polarization types, we find $f_{\text{NL}}^{\text{tll},+} = (11 \pm 8) \times 10^2$ and $f_{\text{NL}}^{\text{tll},-} = (0 \pm 14) \times 10^2$, both of which are consistent with zero at 1.4σ . As such, *we report no detection of equilateral tensor non-Gaussianity, in both the parity-even and parity-odd sectors.* The two constraints are almost entirely uncorrelated; this is as expected, since the parity-odd and parity-even sectors do not mix in an ideal bispectrum estimator [57] and the mask-induced leakage between different parities is seen to be small in Fig. 2. Assuming $f_{\text{NL}}^{\text{tll},+} = f_{\text{NL}}^{\text{tll},-}$ (which is true for a




FIG. 5. Comparison of theoretical tensor-mode non-Gaussianities (blue and orange) to the empirical and theoretical bispectrum variances (green and red), plotting as in Fig. 1. The theoretical curves are computed with $f_{\text{NL}}^{\text{th},\pm} = 1$ and we separate real (blue) and imaginary (orange) contributions. Here, modes generated by parity-even physics are shown on the left (“XYZ+” components), while those from parity-odd physics are on the right (“XYZ-”). We find excellent agreement between simulated (red) and predicted (green) variances across a wide range of scales, with the (projected) ratio plotted in Fig. 1.

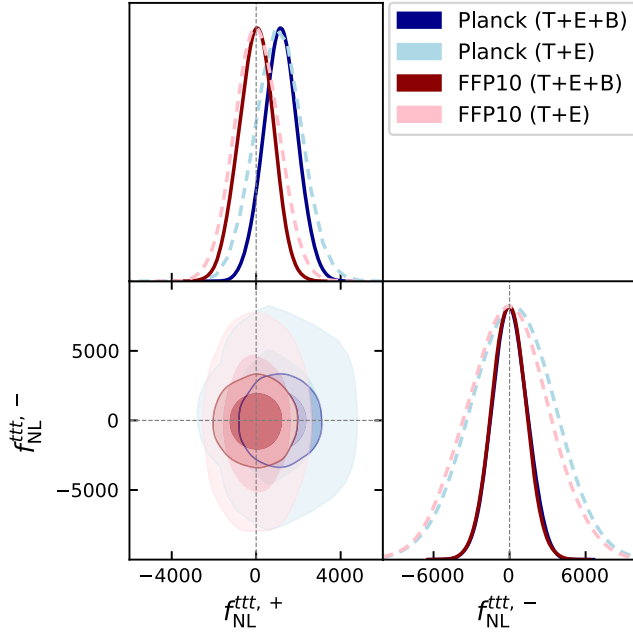


FIG. 6. Constraints on parity-even and parity-odd GW bispectra from the *Planck* PR4 temperature and polarization data set (blue), as well as the mean of 500 FFP10 simulations (red). We show results using both $T + E$ (dashed lines, similar to [42]), and with the inclusion of B -modes (full lines), which is new to this work. B -modes improve the parity-even bounds by $\approx 30\%$ and source the bulk of the parity-odd constraints. We find $f_{\text{NL}}^{\text{th},\pm}$ consistent with zero within the 95% confidence level in all cases. Numerical constraints are given in Tables II and III and all results are obtained by sampling the posterior (23) using emcee [81].

TABLE II. *Planck* PR4 constraints on the equilateral tensor-mode non-Gaussianity parameter, $f_{\text{NL}}^{\text{th}}$. For each analysis, we analyze parity-even (first column) and parity-odd (second column) bispectra separately, and additionally provide a combined constraint (third column), assuming maximal chirality. The fiducial analysis assumes $\ell_{\text{min}} = 2$ and $\ell_{\text{max}} = 500$, and we show the results from *Planck* 2018 [42] in the bottom panel (assuming the SEVEM component-separation pipeline). Most detection significances are small, with the largest signals seen for E -modes or $E + B$ -modes (up to 2.4σ), though this is severely washed out when T -modes are included.

Analysis	$10^{-2} f_{\text{NL}}^{\text{th},+}$	$10^{-2} f_{\text{NL}}^{\text{th},-}$	$10^{-2} f_{\text{NL}}^{\text{th}}$
Fiducial ($T + E + B$)	11 ± 8	-0 ± 14	9 ± 7
T only	-2 ± 20	149 ± 136	5 ± 20
E only	67 ± 30	-384 ± 439	65 ± 30
B only	55 ± 68	-59 ± 59	-10 ± 44
$T + E$	10 ± 11	2 ± 32	10 ± 10
$T + B$	4 ± 14	-8 ± 18	-0 ± 11
$E + B$	40 ± 17	-20 ± 40	37 ± 15
$\ell_{\text{min}} = 4$	13 ± 9	-6 ± 15	8 ± 8
$\ell_{\text{max}} = 375$	12 ± 8	0 ± 14	9 ± 7
No linear term	12 ± 8	-1 ± 14	9 ± 7
T only (<i>Planck</i> 2018)	4 ± 17	90 ± 100	6 ± 16
E only (<i>Planck</i> 2018)	75 ± 75	-790 ± 830	70 ± 75
$T + E$ (<i>Planck</i> 2018)	16 ± 14	2 ± 20	13 ± 12

maximally chiral GW sector), we obtain the combined constraint $f_{\text{NL}}^{\text{ttt}} = (9 \pm 7) \times 10^2$, which is dominated by the parity-even sector.

Excluding B -modes from the analysis is found to inflate the parity-even constraints by $\simeq 30\%$ while increasing the error bar on $f_{\text{NL}}^{\text{ttt},-}$ by a factor $\simeq 2.4$. In the absence of B -modes, parity-odd physics can be probed only through modes with odd $\ell_1 + \ell_2 + \ell_3$ which, by intermediate ℓ , are limited by projection effects. In contrast, TTB , TEB , EEB , and BBB spectra probe parity violation with even $\ell_1 + \ell_2 + \ell_3$, avoiding such a suppression, and sharpening the constraints. If we consider only T -modes, the constraints degrade significantly; this implies that much of the signal-to-noise comes from polarization and its cross-correlation with temperature (noting that the full data set contains ten spectra). A similar story holds for analyses involving just E - or B -modes, with the lack of cross spectra significantly reducing the signal to noise. For E -only or $E + B$ -modes, we find that $f_{\text{NL}}^{\text{ttt},+}$ is nonzero at up to 2.4σ (possibly indicating residual systematic contamination); however, this is strongly inconsistent with the full-data-set result.

To robustly assess the contributions of each field to the overall constraint, we perform a likelihood analysis using each b^{XYZ} bispectrum in turn and plot the fractional contributions to $\sigma^{-2}(f_{\text{NL}}^{\text{ttt},\pm})$ (i.e., the Fisher information) in Fig. 7. For the parity-even sector, we find that constraints are split across a wide variety of channels; $T + E$ -mode correlators make up $\simeq 60\%$ of the Fisher matrix, with the remainder sourced by B -mode spectra (particularly those

with two B -modes, following the above projection arguments). This is quite different to the ratios found for equilateral *scalar* non-Gaussianity in [42], reflecting the utility of spin-2 E -modes in graviton constraints. Parity-odd constraints are instead strongly dominated by the TTB correlator. This matches our expectation, since it is the lowest-noise correlator with even $\ell_1 + \ell_2 + \ell_3$; in this case, the utility of B -modes is abundantly clear, with T -only constraints making up only 1% of the Fisher information.

Finally, it is interesting to ask to what extent these constraints could be obtained from the model-independent analysis of Sec. IV B. In this case one can ask: how large would $f_{\text{NL}}^{\text{ttt},\pm}$ need to be to induce a 1σ deviation in the *Planck* χ^2 value relative to the empirical distribution from FFP10? This is straightforward to compute and leads to the error bars $\sigma(f_{\text{NL}}^{\text{ttt},+}) = 1.3 \times 10^4$, $\sigma(f_{\text{NL}}^{\text{ttt},-}) = 1.9 \times 10^4$ from the fiducial analysis including T -, E -, and B -modes. These constraints are *much* weaker than those given in Table II, reemphasizing our earlier statement that model-specific tests always obtain tighter parameter constraints than blind challenges.

3. Consistency checks

Next, we consider the dependence of our constraints on various hyperparameters (with results given in Table II). First, we consider excising modes with $\ell \leq 4$; this inflates the error bars by $\simeq 10\%$ but does not significantly shift the mean, implying that any bias induced by large-scale residual foregrounds is small. Second, we explore the




FIG. 7. Fractional contribution of each bispectrum to the $f_{\text{NL}}^{\text{ttt},\pm}$ constraints, here parametrized by the Fisher information $\sigma_{XYZ}^{-2}(f_{\text{NL}}^{\text{ttt},\pm})$. We show results for the parity-even and parity-odd sectors in the left and right panels, respectively, and perform all analyses using fiducial hyperparameters, with $\ell \in [2, 500)$. We find that T - and E -modes dominate the parity-even sector (with non-negligible TBB contributions), but the parity-odd sector is primarily influenced by TTB (since it sources spectra with even $\ell_1 + \ell_2 + \ell_3$).

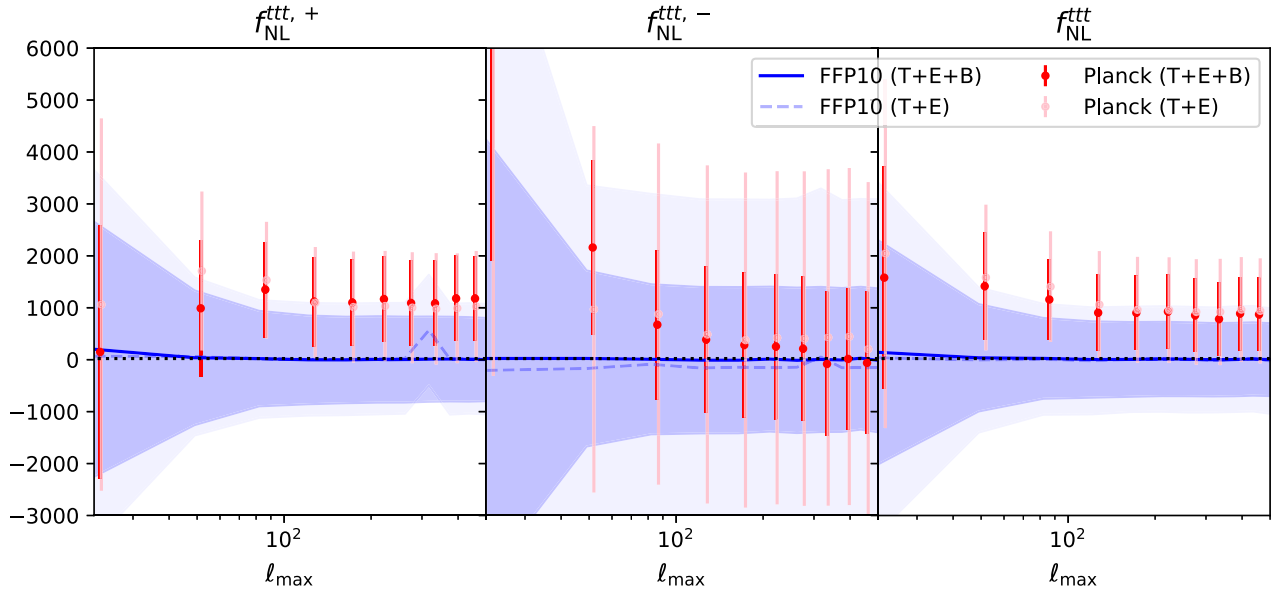


FIG. 8. Constraints on the parity-even (left), parity-odd (center), and combined (right) tensor non-Gaussianity parameters as a function of the maximum scale used in the analysis, ℓ_{\max} . Results from *Planck* PR4 are shown as red data points (or pink, dropping B -modes), while the blue bands indicate results obtained from analyzing the mean of 500 FFP10 simulations. As expected, the constraining power saturates by $\ell_{\max} \simeq 200$.

dependence of the constraints on ℓ_{\max} in Fig. 8. Constraints from both the $T + E$ and $T + E + B$ data sets degrade significantly if we restrict only to very low ℓ ; this occurs due to the form of the tensor transfer functions. Beyond $\ell_{\max} \simeq 200$, the constraints appear to stabilize; indeed, we find negligible changes when the analysis is repeated without the highest-two ℓ bins (Table II). This matches the forecasts of [39,43] and is expected to be similar even for future low-noise experiments, due to the lensing B -mode floor. Different scalings are found for futuristic experiments when invoking delensing, however, as well as for mixed scalar-scalar-tensor [36,43,82,83] and scalar-tensor-tensor [35,82] bispectra. These scalings motivate our choice of ℓ_{\max} (and N_{side}) used throughout this work and indicates that one could perform faster (or more optimally binned) analyses with lower ℓ_{\max} in the future. Furthermore, the bispectrum estimation could be optimized by a factor $\simeq 100$; as seen in Table II, the linear term in the bispectrum estimator makes essentially no difference to the $f_{\text{NL}}^{\text{ttt},\pm}$ constraints and could be dropped.¹³

Last, we have performed extensive verification of our pipeline using the FFP10 simulation suite, leading to the constraints shown in Fig. 6 and Table III. In every analysis, analyzing the mean of 500 simulations returns a result consistent with zero to within 0.2σ . This indicates that the systematic error of our pipeline is under control, at least for

the systematics included in the FFP10 simulation suite, and gives us confidence in the robustness of the above results. We further note that non-Gaussian contributions to the CMB bispectrum, namely, integrated Sachs-Wolfe lensing and polarization-lensing contributions, peak in squeezed regimes and at larger ℓ , and are thus not expected to bias our (large-scale and equilateral) constraints [84–86].

4. Interpretation

The $T + E$ constraints on $f_{\text{NL}}^{\text{ttt},\pm}$ can be directly compared to those of the *Planck* 2018 non-Gaussianity analysis [42] (shown in Table II); the SEVEM pipeline (which is also used in NPIPE) found $f_{\text{NL}}^{\text{ttt},+} = (16 \pm 14) \times 10^2$ with

TABLE III. As in Table II, but for the mean of 500 FFP10 mocks. Results are consistent with zero, indicating that our pipeline is unbiased.

Analysis	$f_{\text{NL}}^{\text{ttt},+}$	$f_{\text{NL}}^{\text{ttt},-}$	$f_{\text{NL}}^{\text{ttt}}$
Fiducial ($T + E + B$)	0 ± 8	-0 ± 14	0 ± 7
T only	1 ± 20	-1 ± 135	1 ± 20
E only	0 ± 30	-20 ± 436	0 ± 30
B only	-0 ± 68	-6 ± 60	-3 ± 44
$T + E$	-0 ± 11	-2 ± 32	-0 ± 10
$T + B$	0 ± 14	-0 ± 18	0 ± 11
$E + B$	-0 ± 17	-2 ± 40	-1 ± 16
$\ell_{\min} = 4$	0 ± 9	-0 ± 15	0 ± 8
$\ell_{\max} = 375$	0 ± 8	-0 ± 14	0 ± 7
No linear term	-1 ± 8	-0 ± 14	0 ± 7

¹³This point is not generic; theoretical models whose signatures are concentrated at the lowest ℓ modes (such as $f_{\text{NL}}^{\text{loc}}$) would strongly benefit from this term.

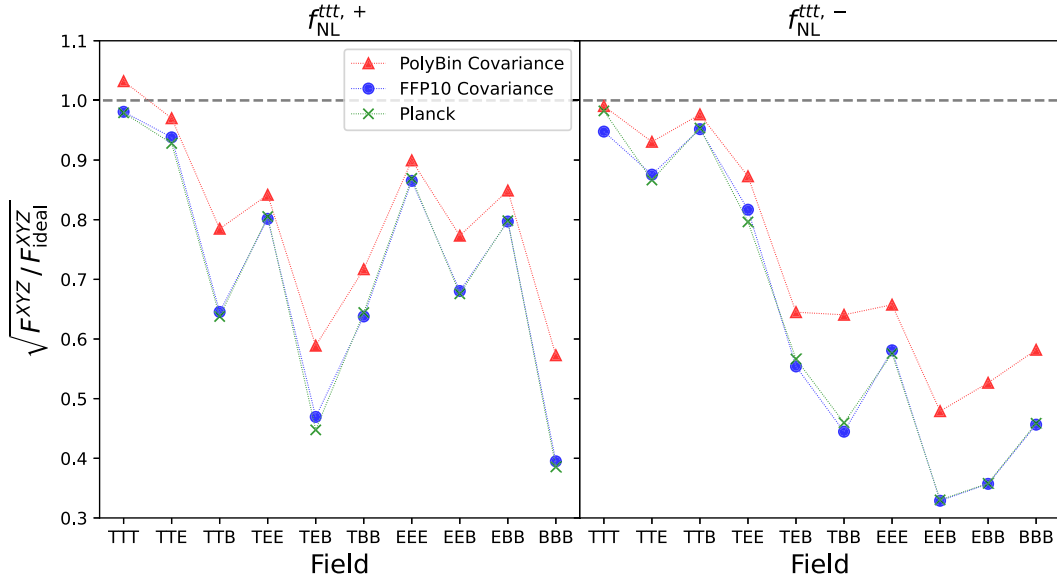


FIG. 9. Ratio of the expected to idealized (inverse) error bars on the non-Gaussianity parameter $f_{\text{NL}}^{\text{ttt},\pm}$ computed from the Fisher matrix of (25) with differing choices of covariance matrix. We show results with the expected covariance including the mask (red), that extracted from 500 FFP10 simulations (blue), and the true error bars obtained from the likelihood analysis of *Planck* (green). We conclude that the nonuniform window function degrades the constraints by $\sim 30\%$, particularly for polarization and parity-odd models, and there is a further $\sim 10\%$ reduction from non-Gaussian contributions to the polarization covariance, encoded within FFP10.

$f_{\text{NL}}^{\text{ttt},-} = (2 \pm 20) \times 10^2$; our parity-even constraint is $\approx 30\%$ tighter, though our parity-odd error bar is degraded by a factor of 1.6. To rationalize this, we note a number of key differences in the analysis. First, our work utilizes an updated *Planck* data set with better treatment of foregrounds and low- ℓ polarization data. Second, we use a lower ℓ_{min} (2 rather than 4, which gives a $\approx 10\%$ improvement). Finally, our analysis utilizes binned bispectra, in contrast to the modal approach [54,55] used in [42]. Our ℓ bins are relatively broad (using only 13 ℓ bins across a wide range of scales), which will lead to some degradation in constraining power. This is not necessarily a limitation of the approach, however; if our analysis was repeated at lower ℓ_{max} (justified by Fig. 8) and without the costly linear term (justified by Table II), one could easily use thinner ℓ bins and thus capture this lost signal to noise; Fisher forecasts using 23 ℓ bins indicate that the constraint on $f_{\text{NL}}^{\text{ttt},+}$ ($f_{\text{NL}}^{\text{ttt},-}$) could be improved by 10–15%. We note, however, that by the inclusion of B -modes, our constraints already represent a significant enhancement over those presented previously, and thus this discussion simply indicates that further improvements are possible.

Finally, it is interesting to compare our constraints to those expected from simplified Fisher forecasts. This is shown in Fig. 9, where we plot the single-field Fisher information,

$$F^{XYZ} = \sum_{\mathbf{b}, \mathbf{b}'} \frac{\partial b^{XYZ}(\mathbf{b})}{\partial f_{\text{NL}}^{\text{ttt},\pm}} [\mathbf{C}^{XYZ}]^{-1}(\mathbf{b}, \mathbf{b}') \frac{\partial b^{XYZ}(\mathbf{b}')}{\partial f_{\text{NL}}^{\text{ttt},\pm}}, \quad (25)$$

for $X, Y, Z \in \{T, E, B\}$ and bins \mathbf{b}, \mathbf{b}' , relative to that expected from an ideal binned analysis with $f_{\text{sky}} = 0.78$ [using the covariance of (19), as shown in Fig. 3]. We consider three Fisher forecasts: (a) using the covariance extracted from PolyBin (inverting the normalization, \mathcal{F}), which fully encapsulates geometric distortions; (b) using the covariance extracted from FFP10 simulations, which further accounts for non-Gaussian contributions to the covariance; and (c) replacing F^{XYZ} with the error bars extracted from the actual analysis of *Planck* data. For TTT bispectra, all ratios are close to unity, indicating that our methodology is optimal (given our binning) and that *Planck* temperature is fully cosmic-variance limited. When polarization is included, the PolyBin constraints degrade by $\approx 10\text{--}40\%$, particularly with the inclusion of B -modes or the parity-odd sector. This indicates that the nonuniform mask significantly complicates our analysis due to its strong scale dependence at low ℓ (given that the expected B -mode signal is small). The FFP10 result (which agrees with *Planck* bounds, as expected for a Gaussian posterior) shows that there is some impact of non-Gaussian covariance (i.e., contributions to the covariance that are not sourced by the two-point function); this reduces the constraining power by $\approx 10\%$. These results are in accordance with our discussion of the covariances themselves in Sec. IV A.

V. DISCUSSION

Primordial GWs represent a unique probe of the early Universe, shedding light on both the inflationary vacuum

and interactions within multifield paradigms. There exist many sources of tensor fluctuations; to distinguish between them, one may consider their statistical properties, in particular their non-Gaussianity and parity. In this work, we have used the full *Planck* temperature and polarization data to place novel constraints on graviton bispectra, considering both parity-even and parity-odd scenarios. Such correlators can be sourced by a number of mechanisms, including nonstandard inflationary vacua, gravitational interactions of gauge fields, scalar-tensor theories of general relativity, nonattractor phases of inflation, Chern-Simons modified gravity, and beyond, and could, in principle be large compared to the two-point function (e.g., [10,11,14,16,23,26–29,35,37]).

Utilizing recently developed binned bispectrum estimators [56–58], we have measured the binned *Planck* T -, E -, and B -mode bispectra across a range of scales, and used these to perform model-agnostic tests of graviton parity violation and non-Gaussianity. Furthermore, we have placed constraints on a popular model of equilateral tensor non-Gaussianity (encoding the gravitational interactions of gauge fields), considering both the parity-even and parity-odd sectors. This work represents the first use of B -modes to probe non-Gaussian inflationary physics; we found that their inclusion enhances constraints on the parity-even sector by $\simeq 30\%$ and dominates parity-odd constraints, particularly through TTB bispectra. Throughout our tests, we found no significant detections; as such, *we report no compelling evidence for GW bispectra, in both the parity-even and parity-odd sectors.*

Comparing to the WMAP, *Planck* 2015, and *Planck* 2018 analysis [40–42], we found significantly enhanced constraints on the tensor non-Gaussianity parameter $f_{\text{NL}}^{\text{t}}_{\text{NL}}$, primarily due to our inclusion of B -modes. Without B -modes, our parity-odd constraints are somewhat worse (though our parity-even constraints improve over the former, due to the lower ℓ_{min}). This is attributed to the wide ℓ bins used in our bispectrum estimator, and could be mollified by reducing the bin width; we forecast that constraints could be improved by 10–15% by doubling the number of linear bins. As noted above, our estimator could be significantly expedited by dropping the linear term (which increases

computation by $\simeq 100\times$ and is of little use for equilateral templates) and by reducing the maximum scale, noting that our constraints saturate by $\ell_{\text{max}} \simeq 200$. We have also compared our results to idealized Fisher forecasts; here, we found that the scale dependence of the mask considerably reduces the precision of B -mode analyses, and there is some non-Gaussian contribution to the covariance, which is well captured by the FFP10 simulations. Such effects will be reduced in the future by better foreground subtraction and understanding of the Galactic mask.

Our constraints on tensor non-Gaussianity will improve considerably in the future. While the large-scale T - and E -modes are quickly becoming cosmic-variance limited, our work has shown that B -modes can significantly enhance both parity-even and parity-odd constraints. Future experiments such as the Simons Observatory, CMB-S4, and LiteBIRD will see such modes measured with much higher precision, and improvements in de-lensing techniques are poised to reduce their remaining cosmic variance, allowing information to also be extracted from smaller scales. For example, the LiteBIRD survey is expected to reach an $f_{\text{NL}}^{\text{t}}_{\text{NL}}$ error bar of $\mathcal{O}(1)$ from (lensed) B -modes alone (representing a $70\times$ improvement over the current constraint) [43]. Furthermore, the techniques considered herein can be simply extended to other scenarios, such as local-type tensor non-Gaussianity or mixed tensor-scalar-scalar non-Gaussianity, which will result in tighter bounds on the corresponding amplitudes. By including the full set of $T + E + B$ -mode correlators, we can strengthen the bounds on a wide variety of primordial tensor phenomena, which will be of paramount importance in future constraints on inflation.

ACKNOWLEDGMENTS

We thank the anonymous referee for insightful comments. O. H. E. P. is a Junior Fellow of the Simons Society of Fellows. M. S. is supported by JSPS KAKENHI Grants No. JP20H05859 and JP23K03390. M. S. also acknowledges the Center for Computational Astrophysics, National Astronomical Observatory of Japan, for providing the computing resources of the Cray XC50.

-
- [1] N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo *et al.* (Planck Collaboration), *Astron. Astrophys.* **641**, A6 (2020).
 - [2] Y. Akrami *et al.* (Planck Collaboration), *Astron. Astrophys.* **641**, A10 (2020).
 - [3] M. Kamionkowski and E. D. Kovetz, *Annu. Rev. Astron. Astrophys.* **54**, 227 (2016).
 - [4] M.-a. Watanabe, S. Kanno, and J. Soda, *Prog. Theor. Phys.* **123**, 1041 (2010).
 - [5] N. Barnaby, J. Moxon, R. Namba, M. Peloso, G. Shiu, and P. Zhou, *Phys. Rev. D* **86**, 103508 (2012).
 - [6] A. Lue, L.-M. Wang, and M. Kamionkowski, *Phys. Rev. Lett.* **83**, 1506 (1999).
 - [7] L. Sorbo, *J. Cosmol. Astropart. Phys.* **06** (2011) 003.

- [8] N. Barnaby, R. Namba, and M. Peloso, *J. Cosmol. Astropart. Phys.* **04** (2011) 009.
- [9] J.L. Cook and L. Sorbo, *J. Cosmol. Astropart. Phys.* **11** (2013) 047.
- [10] R. Namba, M. Peloso, M. Shiraishi, L. Sorbo, and C. Unal, *J. Cosmol. Astropart. Phys.* **01** (2016) 041.
- [11] E. Dimastrogiovanni, M. Fasiello, and T. Fujita, *J. Cosmol. Astropart. Phys.* **01** (2017) 019.
- [12] A. Maleknejad, M. M. Sheikh-Jabbari, and J. Soda, *Phys. Rep.* **528**, 161 (2013).
- [13] E. Komatsu, *Nat. Rev. Phys.* **4**, 452 (2022).
- [14] J.-O. Gong, M. Mylova, and M. Sasaki, *J. High Energy Phys.* **10** (2023) 140.
- [15] S. Kanno and M. Sasaki, *J. High Energy Phys.* **08** (2022) 210.
- [16] O. Ozsoy, M. Mylova, S. Parameswaran, C. Powell, G. Tasinato, and I. Zavala, *J. Cosmol. Astropart. Phys.* **09** (2019) 036.
- [17] X. Gao, T. Kobayashi, M. Yamaguchi, and J. Yokoyama, *Phys. Rev. Lett.* **107**, 211301 (2011).
- [18] X. Gao, T. Kobayashi, M. Shiraishi, M. Yamaguchi, J. Yokoyama, and S. Yokoyama, *Prog. Theor. Exp. Phys.* **2013**, 053E03 (2013).
- [19] V. De Luca, G. Franciolini, A. Kehagias, A. Riotto, and M. Shiraishi, *Phys. Rev. D* **100**, 063535 (2019).
- [20] J. M. Maldacena and G. L. Pimentel, *J. High Energy Phys.* **09** (2011) 045.
- [21] S. Alexander and J. Martin, *Phys. Rev. D* **71**, 063526 (2005).
- [22] J. Soda, H. Kodama, and M. Nozawa, *J. High Energy Phys.* **08** (2011) 067.
- [23] M. Shiraishi, D. Nitta, and S. Yokoyama, *Prog. Theor. Phys.* **126**, 937 (2011).
- [24] N. Bartolo and G. Orlando, *J. Cosmol. Astropart. Phys.* **07** (2017) 034.
- [25] M. C. Guzzetti, N. Bartolo, M. Liguori, and S. Matarrese, *Riv. Nuovo Cimento* **39**, 399 (2016).
- [26] G. Cabass, E. Pajer, D. Stefanyshyn, and J. Supeł, *J. High Energy Phys.* **05** (2022) 077.
- [27] G. Cabass, D. Stefanyshyn, J. Supeł, and A. Thavanesan, *J. High Energy Phys.* **10** (2022) 154.
- [28] L. Bordin and G. Cabass, *J. Cosmol. Astropart. Phys.* **07** (2020) 014.
- [29] G. Cabass, *J. Cosmol. Astropart. Phys.* **12** (2021) 001.
- [30] E. Pajer, *J. Cosmol. Astropart. Phys.* **01** (2021) 023.
- [31] J. M. Maldacena, *J. High Energy Phys.* **05** (2003) 013.
- [32] A. Agrawal, T. Fujita, and E. Komatsu, *Phys. Rev. D* **97**, 103526 (2018).
- [33] M. Shiraishi, C. Hikage, R. Namba, T. Namikawa, and M. Hazumi, *Phys. Rev. D* **94**, 043506 (2016).
- [34] B. Thorne, T. Fujita, M. Hazumi, N. Katayama, E. Komatsu, and M. Shiraishi, *Phys. Rev. D* **97**, 043506 (2018).
- [35] N. Bartolo, G. Orlando, and M. Shiraishi, *J. Cosmol. Astropart. Phys.* **01** (2019) 050.
- [36] A. J. Duivenvoorden, P. D. Meerburg, and K. Freese, *Phys. Rev. D* **102**, 023521 (2020).
- [37] M. Mylova, [arXiv:1912.00800](https://arxiv.org/abs/1912.00800).
- [38] M. Shiraishi, *J. Cosmol. Astropart. Phys.* **06** (2012) 015.
- [39] M. Shiraishi, A. Ricciardone, and S. Saga, *J. Cosmol. Astropart. Phys.* **11** (2013) 051.
- [40] M. Shiraishi, M. Liguori, and J. R. Fergusson, *J. Cosmol. Astropart. Phys.* **01** (2015) 007.
- [41] P. A. R. Ade *et al.* (Planck Collaboration), *Astron. Astrophys.* **594**, A17 (2016).
- [42] Y. Akrami *et al.* (Planck Collaboration), *Astron. Astrophys.* **641**, A9 (2020).
- [43] M. Shiraishi, *Front. Astron. Space Sci.* **6**, 49 (2019).
- [44] M. Shiraishi, *Phys. Rev. D* **94**, 083503 (2016).
- [45] O. H. E. Philcox, J. Hou, and Z. Slepian, [arXiv:2108.01670](https://arxiv.org/abs/2108.01670).
- [46] R. N. Cahn, Z. Slepian, and J. Hou, *Phys. Rev. Lett.* **130**, 201002 (2023).
- [47] J. Hou, Z. Slepian, and R. N. Cahn, *Mon. Not. R. Astron. Soc.* **522**, 5701 (2023).
- [48] O. H. E. Philcox, *Phys. Rev. D* **106**, 063501 (2022).
- [49] G. Cabass, M. M. Ivanov, and O. H. E. Philcox, *Phys. Rev. D* **107**, 023523 (2023).
- [50] C. Creque-Sarbinowski, S. Alexander, M. Kamionkowski, and O. Philcox, *J. Cosmol. Astropart. Phys.* **11** (2023) 029.
- [51] O. H. E. Philcox, *Phys. Rev. Lett.* **131**, 181001 (2023).
- [52] O. H. E. Philcox and M. Shiraishi, [arXiv:2308.03831](https://arxiv.org/abs/2308.03831).
- [53] W. R. Coulton and D. N. Spergel, *J. Cosmol. Astropart. Phys.* **10** (2019) 056.
- [54] M. Shiraishi, M. Liguori, and J. R. Fergusson, *J. Cosmol. Astropart. Phys.* **05** (2014) 008.
- [55] M. Shiraishi, M. Liguori, J. R. Fergusson, and E. P. S. Shellard, *J. Cosmol. Astropart. Phys.* **06** (2019) 046.
- [56] O. H. E. Philcox, *Phys. Rev. D* **107**, 123516 (2023).
- [57] O. H. E. Philcox, *Phys. Rev. D* **108**, 063506 (2023).
- [58] O. H. E. Philcox, PolyBin: Binned polyspectrum estimation on the full sky, Astrophysics Source Code Library, record ascl:2307.020 (2023).
- [59] M. Kamionkowski and T. Souradeep, *Phys. Rev. D* **83**, 027301 (2011).
- [60] M. Bucher, B. Racine, and B. van Tent, *J. Cosmol. Astropart. Phys.* **05** (2016) 055.
- [61] M. Bucher, B. Van Tent, and C. S. Carvalho, *Mon. Not. R. Astron. Soc.* **407**, 2193 (2010).
- [62] K. M. Smith and M. Zaldarriaga, *Mon. Not. R. Astron. Soc.* **417**, 2 (2011).
- [63] K. M. Smith, L. Senatore, and M. Zaldarriaga, [arXiv:1502.00635](https://arxiv.org/abs/1502.00635).
- [64] M. Shiraishi, D. Nitta, S. Yokoyama, K. Ichiki, and K. Takahashi, *Prog. Theor. Phys.* **125**, 795 (2011).
- [65] M. Shiraishi, S. Yokoyama, K. Ichiki, and K. Takahashi, *Phys. Rev. D* **82**, 103505 (2010).
- [66] M. Shiraishi, *Mod. Phys. Lett. A* **31**, 1640003 (2016).
- [67] Y. Watanabe and E. Komatsu, [arXiv:2004.04350](https://arxiv.org/abs/2004.04350).
- [68] L. Mirzaghali, E. Komatsu, K. D. Lozanov, and Y. Watanabe, *J. Cosmol. Astropart. Phys.* **06** (2020) 024.
- [69] N. Bartolo, S. Matarrese, M. Peloso, and M. Shiraishi, *J. Cosmol. Astropart. Phys.* **07** (2015) 039.
- [70] Y. Akrami *et al.* (Planck Collaboration), *Astron. Astrophys.* **643**, A42 (2020).
- [71] E. Rosenberg, S. Gratton, and G. Efstathiou, *Mon. Not. R. Astron. Soc.* **517**, 4620 (2022).
- [72] M. Tristram *et al.*, *Astron. Astrophys.* **647**, A128 (2021).
- [73] Y. Akrami *et al.* (Planck Collaboration), *Astron. Astrophys.* **641**, A4 (2020).

- [74] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelman, *Astrophys. J.* **622**, 759 (2005).
- [75] H. F. Gruetjen, J. R. Fergusson, M. Liguori, and E. P. S. Shellard, *Phys. Rev. D* **95**, 043532 (2017).
- [76] O. H. E. Philcox, *Phys. Rev. D* **103**, 103504 (2021).
- [77] O. H. E. Philcox, *Phys. Rev. D* **104**, 123529 (2021).
- [78] A. J. S. Hamilton and M. Tegmark, *Mon. Not. R. Astron. Soc.* **312**, 285 (2000).
- [79] M. Reinecke and D. S. Seljebotn, *Astron. Astrophys.* **554**, A112 (2013).
- [80] W. R. Coulton, O. H. E. Philcox, and F. Villaescusa-Navarro, *Phys. Rev. D* **109**, 023531 (2024).
- [81] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, *Publ. Astron. Soc. Pac.* **125**, 306 (2013).
- [82] M. Shiraishi, *J. Cosmol. Astropart. Phys.* 11 (2013) 006.
- [83] G. Domènech, T. Hiramatsu, C. Lin, M. Sasaki, M. Shiraishi, and Y. Wang, *J. Cosmol. Astropart. Phys.* 05 (2017) 034.
- [84] W. Hu, *Phys. Rev. D* **62**, 043007 (2000).
- [85] A. Lewis, A. Challinor, and D. Hanson, *J. Cosmol. Astropart. Phys.* 03 (2011) 018.
- [86] R. Pearson, A. Lewis, and D. Regan, *J. Cosmol. Astropart. Phys.* 03 (2012) 011.