

What can be learned from the lensed cosmic microwave background B -mode polarization power spectrum?

Sarah Smith and Anthony Challinor

Astrophysics Group, Cavendish Laboratory, J J Thomson Avenue, Cambridge CB3 0HE, United Kingdom

Graça Rocha

Astrophysics Group, Cavendish Laboratory, J J Thomson Avenue, Cambridge CB3 0HE, United Kingdom

Observational Cosmology, California Institute of Technology, Pasadena California 91125, USA

Jet Propulsion Laboratory, Pasadena California 91109, USA

(Received 24 November 2005; published 25 January 2006)

The effect of weak gravitational lensing on the cosmic microwave background (CMB) temperature anisotropies and polarization will provide access to cosmological information that cannot be obtained from the primary anisotropies alone. We compare the information content of the lensed B -mode polarization power spectrum, properly accounting for the non-Gaussian correlations between the power on different scales, with that of the unlensed CMB fields and the lensing potential. The latter represent the products of an (idealized) optimal analysis that exploits the lens-induced non-Gaussianity to reconstruct the fields. Compressing the non-Gaussian lensed CMB into power spectra is wasteful and leaves a tight degeneracy between the equation of state of dark energy and neutrino mass that is much stronger than in the more optimal analysis. Despite this, a power-spectrum analysis will be a useful first step in analyzing future B -mode polarization data. For this reason, we also consider how to extract accurate parameter constraints from the lensed B -mode power spectrum. We show with simulations that for cosmic-variance-limited measurements of the lensed B -mode power, including the non-Gaussian correlations in existing likelihood approximations gives biased parameter results. We develop a more refined likelihood approximation that performs significantly better. This new approximation should also be of more general interest in the wider context of parameter estimation from Gaussian CMB data.

DOI: [10.1103/PhysRevD.73.023517](https://doi.org/10.1103/PhysRevD.73.023517)

PACS numbers: 98.80.Es, 98.70.Vc

I. INTRODUCTION

The goal of many current and forthcoming cosmic microwave background (CMB) experiments is to measure the polarization of the CMB with increasing accuracy (e.g. QUaD [1], BICEP [2], EBEX [3], QUIET¹ and Clover [4]). There are several motivations for this: polarization measurements are an important test of the consistency of the cosmological model, as well containing new information, and can break parameter degeneracies which are present if only the CMB temperature power spectrum is analyzed [5]. A key motivation is the fact that the curl-like B -mode polarization is not generated directly from scalar perturbations at last scattering, and thus can potentially reveal the presence of a primordial gravitational-wave background, if it exists at a high enough level [6,7]. A detection of primordial large-angle B -mode polarization would thus provide a measure of the expansion rate and hence energy scale of inflation. Apart from the reionization bump around $\ell \approx 5$, which depends on the optical depth to last scattering, the gravitational-wave signal peaks at multipoles $\ell \sim 100$, corresponding to the angle subtended by the horizon size at last scattering, and so large area surveys (a few hundred deg²) will be required to detect it. On smaller scales, the dominant contribution to B -mode po-

larization is expected to be the weak lensing of E -modes by large-scale structure along the line of sight [8].

The lensing signal is sensitive to the gravitational potential, and hence the clustering of matter, mostly at redshifts $\lesssim 10$ [9]. (At the peak of the lensing deflection power spectrum at multipoles $\ell \sim 40$, more than 95% of the power arises from $z < 10$.) Lensing is therefore sensitive to parameters that have little direct effect on the pattern of primordial fluctuations, but affect the late-time growth of large-scale structure. The lensing signal can thus provide us with information which would otherwise be absent from the CMB [10,11], such as the properties of dark energy and sub-eV neutrino masses [12]. However, the lensed B -modes are significantly non-Gaussian, and this has been shown to degrade markedly the amount of information present in the B -mode power spectrum [13]. The non-Gaussianity introduced by weak lensing is also present to a much lesser extent in the CMB temperature anisotropies and E -mode polarization signal. It has recently been shown that neglecting this non-Gaussianity does not significantly bias parameter constraints from Planck-quality data, although the effect of weak lensing on the power spectra cannot be neglected [14].

In the present paper we consider more fully the question of what parameter information is contained in the lensed B -mode power spectrum and how to analyze future spectral data reliably. We concentrate on the B -mode spectrum here for the following reason. The expectation is that for

¹<http://quiet.uchicago.edu/>

future experiments with sufficient signal-to-noise to image the lens-induced B -modes, most of the additional parameter information from the lensing effect on the CMB power spectra will come from the B -modes. Although the information content of the B -mode power spectrum is relatively more affected by the non-Gaussianity than the temperature (T) and E -mode polarization, that latter must contend with the cosmic variance of the dominant primary (unlensed) contributions. It is well known that for Gaussian fields the observed power spectrum is a sufficient statistic for parameter estimation under ideal survey conditions, but that this is not the case for non-Gaussian fields such as the lensed CMB. Compressing the lensed CMB fields to their power spectra is wasteful leading to a loss of cosmological information. A more optimal analysis would be to use the non-Gaussianity to reconstruct an estimate of the lensing deflection field and the unlensed CMB fields [15,16], or to work directly from the correct non-Gaussian likelihood function [17,18]. While such analyses are a worthy goal to strive for, they will likely be difficult to implement in practice in the presence of real-world complications such as inhomogeneous noise, complex survey geometries and foreground residuals. A simpler, and probably more robust method to deal with near-future lensed data is to work directly with the lensed fields in a conventional power-spectrum analysis.

This paper is organized as follows. In Sec. II we review how B -mode polarization is generated by the weak lensing of E -modes, and consider the cosmological parameters which can be constrained using this information, and the degeneracies between them. In Sec. III we estimate these parameters from simulated lensed data under the assumption that the data are Gaussian, in order to show directly that this leads to false conclusions. In order to account for the non-Gaussianity we need to understand the correlation between the measured power on different scales, and we calculate this in Sec. IV following Ref. [13]. In Sec. V we incorporate the non-Gaussian covariance in existing likelihoods used for the analysis of (Gaussian) CMB data but show that this too leads to biased results in simulations. The reason for this deficiency appears to arise from inaccuracies in the current likelihood approximations when applied to data with few degrees of freedom (which extends to higher multipoles for the lensed B -modes than for Gaussian fields). We introduce a more accurate likelihood function in Sec. VI which we show to perform much better on simulated lensed data. An appendix provides further details of the derivation of this new likelihood.

II. COSMOLOGICAL PARAMETERS

The action of lensing on the part of the temperature anisotropy and polarization that is generated at recombination is a simple remapping that in the flat-sky approximation can be written as




FIG. 1 (color online). Unlensed $\tilde{E}\tilde{E}$ power spectrum (red solid line) and power spectrum of lensing potential (blue dotted line) which together generate the lensed BB power spectrum (shown by green dashed line, multiplied by a factor of 400).

$$X(\mathbf{n}) = \tilde{X}(\mathbf{n} + \nabla\phi). \quad (1)$$

Here $\tilde{X}(\mathbf{n})$ is the unlensed variable (T , or the Stokes parameters Q and U), $X(\mathbf{n})$ the lensed value and ϕ the lensing deflection field. On the sphere, we need to take care to interpret correctly the meaning of $\nabla\phi$ [19] and to take into account the position-dependence of the co-ordinate basis for polarization. The root-mean-square lensing deflection angle ~ 3 arcmin and so the lensing action on CMB fields well above this scale can be evaluated with the gradient approximation:

$$X(\mathbf{n}) \approx \tilde{X}(\mathbf{n}) + \nabla\tilde{X} \cdot \nabla\phi. \quad (2)$$

Within this approximation, and in the flat-sky limit, we can express the lensed B -mode power spectrum as a convolution of the unlensed E -mode power spectrum with the power spectrum of the lensing potential ϕ [13]:

$$C_{\ell}^{BB} = \int \frac{d^2\ell'}{(2\pi)^2} C_{\ell'}^{\tilde{E}\tilde{E}} C_{|\ell-\ell'|}^{\phi\phi} W^2(\ell, \ell'), \quad (3)$$

where function W is given by

$$W(\ell, \ell') = \ell' \cdot (\ell - \ell') \sin 2(\phi_{\ell} - \phi_{\ell'}), \quad (4)$$

with $\ell = (\ell \cos\phi_{\ell}, \ell \sin\phi_{\ell})$. This is illustrated in Fig. 1, which shows the unlensed E -mode and lensing potential power spectra, together with the resulting lensed B -mode power spectrum. The power spectra were calculated using CAMB² which includes corrections for curved sky effects and the breakdown of the gradient approximation [20].

The lensed B -mode power spectrum is thus affected by cosmological parameters which affect either the primary

²<http://camb.info/>

(unlensed) CMB power spectrum $C_{\ell}^{\bar{E}\bar{E}}$ or the lensing power $C_{\ell}^{\phi\phi}$. Many of these parameters can be well constrained by the temperature and E -mode power spectra, and the lens-induced B -modes do not contribute significantly to improving these constraints. Here we concentrate on those parameters that are degenerate with respect to the unlensed CMB spectra. As is well known, the unlensed spectra can provide tight constraints on the physical densities in baryons and cold dark matter, the angular diameter distance to last scattering, the primordial power spectrum, and the optical depth to reionization [21]. Neutrinos with masses below $\sim 0.3\text{eV}$, as implied by current analyses of large-scale structure and the Ly- α forest [22], are relativistic at recombination and the only effect of their mass on the unlensed CMB is via the angular diameter distance, d_A , and a small large-scale contribution to the late integrated Sachs-Wolfe (ISW) effect [23]. Similarly, dark-energy parameters, such as the current energy density $\Omega_{\Lambda}h^2$, and equation of state w are felt only through d_A and the ISW effect. In the inflation-inspired flat models considered here, the parameters $\Omega_{\Lambda}h^2$, w , and the current density in massive neutrinos $\Omega_{\nu}h^2$ (or equivalently the neutrino masses) are highly degenerate with respect to the unlensed CMB, even for cosmic-variance-limited observations. They are collectively constrained only by d_A , which is now accurately determined to be 13.7 ± 0.5 Gpc from the CMB alone [24]. Adding curvature or evolution in the dark-energy sector increases the dimensionality of the degeneracy.

While this ‘‘geometric’’ degeneracy can be easily broken with the inclusion of external data [21], it is still worthwhile to consider to what extent the lensing of the CMB helps given its relatively simple and well-understood physics. An early analysis of the temperature and E -mode polarization power spectra showed that the weak lensing effect on the power spectra can go some way to breaking the geometric degeneracy [11] since the late-time evolution of the gravitational potential responds differently to variations in the parameters that are geometrically degenerate. Since for multipoles $\ell > 100$ the B -mode power spectrum is expected to be dominated by weak lensing, this is potentially a much more sensitive probe of the lensing effect than the temperature or E -mode power since there is no cosmic variance coming from unlensed B -mode polarization. For a given survey, once the sensitivity reaches the limit for imaging the lens-induced B -modes (better than $5\mu\text{K-arcmin}$) we can expect these to be most constraining. To reach the ultimate cosmic-variance limit considered later in this paper, we require sensitivities around $1\mu\text{K-arcmin}$ over a large fraction of the sky. This is similar to the specifications being discussed for a post-Planck orbital experiment dedicated to CMB polarization³.

In Fig. 2 we show the derivatives of the lensing potential and B -mode power spectra with respect to w and $\Omega_{\nu}h^2$ in a




FIG. 2 (color online). Derivatives of the lensing power spectrum $\ell^4 C_{\ell}^{\phi\phi}$ (top) and the lensed B -mode power spectrum $\ell(\ell + 1)C_{\ell}^{BB}/(2\pi)$ (bottom), keeping θ_D fixed (see text). Red solid line: derivatives with respect to $\Omega_{\nu}h^2$; blue dashed line: derivatives with respect to w (amplified by a factor of 350 in both plots).

flat universe with the physical baryon and cold dark matter densities held fixed at $\Omega_b h^2 = 0.024$ and $\Omega_c h^2 = 0.111$. The derivatives are taken around $w = -1$ (a cosmological constant) and $\Omega_{\nu}h^2 = 0.004$ with the ratio θ_D of the sound horizon at last scattering to the angular diameter distance held fixed. If we parametrize flat models by θ_D , w and $\Omega_{\nu}h^2$, both the Hubble parameter h and $\Omega_{\Lambda}h^2$ are derived parameters. Our choice of parametrization is motivated by the geometric degeneracy: the unlensed CMB accurately determines all our parameters except w and $\Omega_{\nu}h^2$ which are themselves very poorly constrained in the absence of external information. The power spectra used to compute these derivatives were computed with the input unlensed power spectra set to zero above the maximum multipole value $\ell_{\text{max}} = 2048$ to be consistent with the analyses of simulations presented in the following sections of this paper. We have assumed three families of neutrinos with equal masses. Since $\sum m_{\nu} = \Omega_{\nu}h^2 \times 94$ eV when all fam-

³<http://universe.nasa.gov/program/inflation.html>

ilies are nonrelativistic⁴, the mass in the fiducial model is 0.13 eV. The total mass is close to the current best limits from cosmological probes [22]. Given the measurements of (squared) mass differences from atmospheric and solar neutrino experiments, our value of $\Omega_\nu h^2$ is just at the limit where we can assume mass degeneracy [25]. We note also that we have not included any nonlinear corrections to the matter power spectrum in computing the spectra in Fig. 2 since the HALOFIT [26] fitting employed in CAMB does not currently support massive neutrinos or models with $w \neq -1$. Nonlinearities increase the lensing power spectrum beyond $l \sim 300$, and have an effect $>5\%$ on all scales for the B -mode spectrum, and a much larger effect on small scales [20]. Nonlinearities should thus properly be included in any future data analysis.

The effect of changes in w and $\Omega_\nu h^2$ on the lensing power spectrum has been discussed in Ref. [27]. The main effect of changes of w is through the change in the expansion rate. An increase in w at fixed θ_D (or, equivalently, d_A) requires a reduction in the current dark-energy density and hence h . However, increasing w causes dark energy to dominate earlier and the net effect is an initial enhancement of the expansion rate over that for $w = -1$; this causes the gravitational potential to decay earlier and suppresses the lensing effect. Hence the derivative of the lensing potential power spectrum with respect to w is negative; see Fig. 2. The effect on the potential is almost independent of scale, but its time dependence causes a larger fractional change in lensing power on large scales. If we fix the current physical density in dark energy, the effect of massive neutrinos on the gravitational potential vanishes on large scales (larger than the Jeans length of the neutrinos when they first become nonrelativistic) since the increased expansion rate is offset by the neutrino clustering. The effect on the lensing power spectrum thus also vanishes at large ℓ . On scales below the neutrino Jeans length today the neutrinos have never clustered and their mass gives a scale-independent suppression of the gravitational potential similar to the effect of increasing w . On intermediate scales, there is a scale-dependent suppression of the lensing power. The additional effect of keeping θ_D fixed is to subtract off a small proportion of the suppression that arises when $\Omega_\Lambda h^2$ is increased so that neutrino mass then gives a small positive enhancement of lensing power at low ℓ and a suppression for larger ℓ (Fig. 2). The different scale dependence of the effects of neutrino mass and w allows them to be separated if the lensing potential

⁴We note that $\Omega_\nu h^2$ tends to 1.7×10^{-5} in the limit that the neutrino masses tend to zero (for three families of leptons). It is a good approximation to take the neutrino mass of a family to be proportional to its contribution to $\Omega_\nu h^2$ provided that $m_\nu > 0.004$ eV. Since in the November 2004 release of CAMB that we use the (degenerate) masses are calculated as $\sum m_\nu = \Omega_\nu h^2 \times 94$ eV, it is important to avoid calculating derivatives with respect to $\Omega_\nu h^2$ around $\Omega_\nu h^2 = 0$.

can be accurately determined. In this way one can determine *both* w and $\Omega_\nu h^2$ accurately, whereas neither (nor any combination) can be determined accurately from the unlensed CMB alone. Forecasts for such constraints were given in Ref. [27], which used for the errors on $C_l^{\phi\phi}$ those expected from an application of the quadratic reconstruction technique of Ref. [16]. This exploits the local scale-scale correlations in the non-Gaussian lensed CMB fields to reconstruct a (noisy) estimate of the deflection field. Improvements in the reconstruction of the deflection field may be possible with more optimal techniques [18]. Reference [27] found that a post-Planck experiment could determine w to within an error of 0.18 and detect the mass of a single family of massive neutrinos if $m_\nu > 0.04$ eV.

Figure 2 shows that w and $\Omega_\nu h^2$ are largely degenerate with respect to their effect on the B -mode power spectrum. This is because for no ℓ is the dominant contribution to C_ℓ^{BB} coming from the largest angle lenses for which there is a clear difference between the effects of varying w and $\Omega_\nu h^2$; for $\ell \gtrsim 200$, the effect of varying w or $\Omega_\nu h^2$ is almost degenerate in the lensing potential power spectrum and hence in C_ℓ^{BB} . In Fig. 3 we plot the derivatives $C_{\ell'}^{\phi\phi} dC_\ell^{BB}/dC_{\ell'}^{\phi\phi}$ as a function of ℓ' for several values of ℓ . For any ℓ , the integral of the kernel over ℓ' gives the lensed B -mode power spectrum at that ℓ (in the gradient approximation, Eq. (1)). Similarly, for the parameter variations considered here that leave the unlensed E -mode power spectrum unchanged, the convolution of the kernel $C_{\ell'}^{\phi\phi} dC_\ell^{BB}/dC_{\ell'}^{\phi\phi}$ with the fractional change in $C_{\ell'}^{\phi\phi}$ gives the change in the lensed B -mode spectrum. The shape of the kernels in Fig. 3 can be understood as follows. For $\ell \leq 200$, the dominant contribution to the integral in Eq. (3) comes from $\ell' \gg \ell$ since the E -modes have very little power at the large-scale ℓ . In this limit, the lensed




FIG. 3 (color online). Kernels $C_{\ell'}^{\phi\phi} dC_\ell^{BB}/dC_{\ell'}^{\phi\phi}$ for $\ell = 10$ (red), $\ell = 500$ (blue) and $\ell = 1000$ (green). These give (approximately) the contribution to the lensed B -mode spectrum from lenses at scale ℓ' .

B -mode spectrum is approximately white and the kernel reduces to

$$C_{\ell'}^{\phi\phi} \frac{dC_{\ell}^{BB}}{dC_{\ell'}^{\phi\phi}} \approx \frac{1}{4\pi} \ell'^5 C_{\ell'}^{\phi\phi} C_{\ell'}^{\tilde{E}\tilde{E}}, \quad (5)$$

for $\ell \ll \ell'$. For ℓ close to 1000, around the peak of the E -mode spectrum, the dominant contribution to C_{ℓ}^{BB} is from larger scale lenses and the kernel is of the form

$$C_{\ell'}^{\phi\phi} \frac{dC_{\ell}^{BB}}{dC_{\ell'}^{\phi\phi}} \approx \frac{1}{4\pi} \ell'^5 C_{\ell'}^{\phi\phi} \langle C_{\ell}^{\tilde{E}\tilde{E}} \rangle_{\ell'}, \quad (6)$$

for $\ell > \ell'$. Here, $\langle C_{\ell}^{\tilde{E}\tilde{E}} \rangle_{\ell'}$ is a smoothed version of $C_{\ell}^{\tilde{E}\tilde{E}}$ where the smoothing is with a bimodal kernel of total width $2\ell'$. If $\ell' \ll 200$ the smoothing has little effect and $\langle C_{\ell}^{\tilde{E}\tilde{E}} \rangle_{\ell'} \approx C_{\ell}^{\tilde{E}\tilde{E}}$. Finally, for $\ell \gtrsim 5000$ (not shown in Fig. 2), the B -mode power arises from lenses at the same scale across which the CMB may be approximated by a gradient [28]. In this limit, the kernel is roughly

$$C_{\ell'}^{\phi\phi} \frac{dC_{\ell}^{BB}}{dC_{\ell'}^{\phi\phi}} \sim \frac{1}{4} \ell'^2 C_{\ell'}^{\phi\phi} \left(\frac{1}{2\pi} \int \frac{dL}{L} L^4 C_L^{\tilde{E}\tilde{E}} \right) \delta_{\ell\ell'}, \quad (7)$$

where the factor in brackets is the mean-squared gradient of the polarization.

Since w and $\Omega_{\nu} h^2$ are essentially degenerate in the B -mode power spectrum, the consequence of this is that, if we use only the lensed B -mode power spectrum and do not consider higher-order statistics, there is a tight degeneracy between the two parameters. This is still an improvement over what can be achieved from the unlensed CMB spectra alone which give no constraints on either parameter, but clearly falls short of what would be achieved if the lensing potential could be reconstructed directly. While this is a clear motivation for developing further such reconstruction techniques, for the reasons mentioned in Sec. I, it will still be worthwhile to perform an initial power-spectrum-based analysis. Furthermore, external data on e.g. the Hubble parameter, will assist in breaking the degeneracy between w and $\Omega_{\nu} h^2$ in a power-spectrum-only analysis. For this reason, we now consider in more detail how to analyze future B -mode power-spectrum data taking careful account of the dependencies between power on different scales.

III. ANALYSIS ASSUMING GAUSSIANTY

We begin by illustrating the effect that non-Gaussianity has on estimated parameter constraints by (wrongly) analyzing simulated lensed maps with the likelihood function appropriate for Gaussian fields. This extends the work of Ref. [13] who showed that the Fisher estimate of the error on the amplitude of the lensed B -mode spectrum is underestimated if Gaussianity is assumed.

We used the publicly-available LensPix⁵ code [14] (which is based on a modified version of HEALPix 1.2⁶) to create simulations of the lensed CMB on the full sky. So that the effects of non-Gaussianity could be clearly seen we considered the idealised full-sky case with no noise or foreground sources. The lensed map was calculated directly at the pixel centers of the deflection map from the spherical multipoles of the unlensed polarization. This avoids introducing any unwanted bias due to the effects of pixelization error which may arise when using the alternative interpolation method in the LensPix code, but is much slower and limited the number of simulations that we could produce. The sky was generated using the HEALPix resolution parameter $N_{\text{side}} = 1024$ (giving 3.4-arcmin pixels) with $\ell_{\text{max}} = 2048$ to avoid aliasing. The theoretical power spectra used in the analysis were calculated also using $\ell_{\text{max}} = 2048$. This leads to a lensed B -mode power spectrum that lacks power particularly on smaller scales, but our analysis is self-consistent.

We restricted the analysis to constraining only the values of w and the tensor/scalar ratio r , within a flat dark-energy CDM cosmology with no massive neutrinos. The strong degeneracy between w and $\Omega_{\nu} h^2$ means that it was not useful to vary both parameters, and omitting massive neutrinos significantly speeds up the computation of the theoretical power spectra. The models used had scale-invariant curvature fluctuations ($n_s = 1$) with no running. The spectral index of tensor fluctuations was set to $n_t = -r/8$, in accordance with predictions of slow-roll inflation (see, e.g. Ref. [29] and references therein). The values of the physical baryon and CDM densities, $\Omega_b h^2$ and $\Omega_c h^2$, were held fixed along with θ_D , so that Ω_{Λ} and h are derived parameters. Two fiducial models were used, both with $w = -1$, but the second with a much lower value of r to see if the non-Gaussianity had a significant effect on the estimated values of r . Model A had a high value of $r = 0.4$ (close to current upper limits [22]) with $\Omega_b h^2 = 0.024$, $\Omega_c h^2 = 0.111$, $H_0 = 0.66 \text{ kms}^{-1} \text{ Mpc}^{-1}$, the amplitude of scalar fluctuations on 0.05 Mpc^{-1} scales $A_s = 2.6 \times 10^{-9}$ and optical depth $\tau = 0.2$. Model B used $r = 0.01$, $\Omega_b h^2 = 0.0234$, $\Omega_c h^2 = 0.111$, $H_0 = 73 \text{ kms}^{-1} \text{ Mpc}^{-1}$, $A_s = 2.3 \times 10^{-9}$ and $\tau = 0.14$. (The value of $A_s e^{-2\tau}$ was the same for both models.)

We extracted the measured power spectrum \hat{C}_{ℓ}^{BB} of the simulations and constrained r and w with the likelihood appropriate to full-sky, noise-free observations of Gaussian fields:

$$P(\hat{C}|C) \propto \prod_{\ell} \hat{C}_{\ell}^{\ell-1/2} C_{\ell}^{-(\ell+1/2)} \exp\left\{-\frac{(2\ell+1)\hat{C}_{\ell}}{2C_{\ell}}\right\}, \quad (8)$$

where C_{ℓ} is the theoretical power spectrum. The normalization of the likelihood function with respect to \hat{C}_{ℓ} is

⁵<http://cosmologist.info/lenspix>

⁶<http://healpix.jpl.nasa.gov>

usually dropped since in practice it is a fixed quantity when constraining parameters. We assumed flat priors on the values of w and r .

Figure 4 shows the parameter constraints obtained from three Model-A simulations and seven Model-B simulations using the likelihood in Eq. (8). It can be seen that the constraints are, in most cases, inconsistent with the input value of $w = -1$. However, the marginalized constraints on r are in accordance with the input values. There appears to be a bias towards low values of w , and the spread of the maximum-likelihood values indicates that the constraints on w are too tight. The latter is consistent with the findings of the Fisher analysis in Ref. [13].

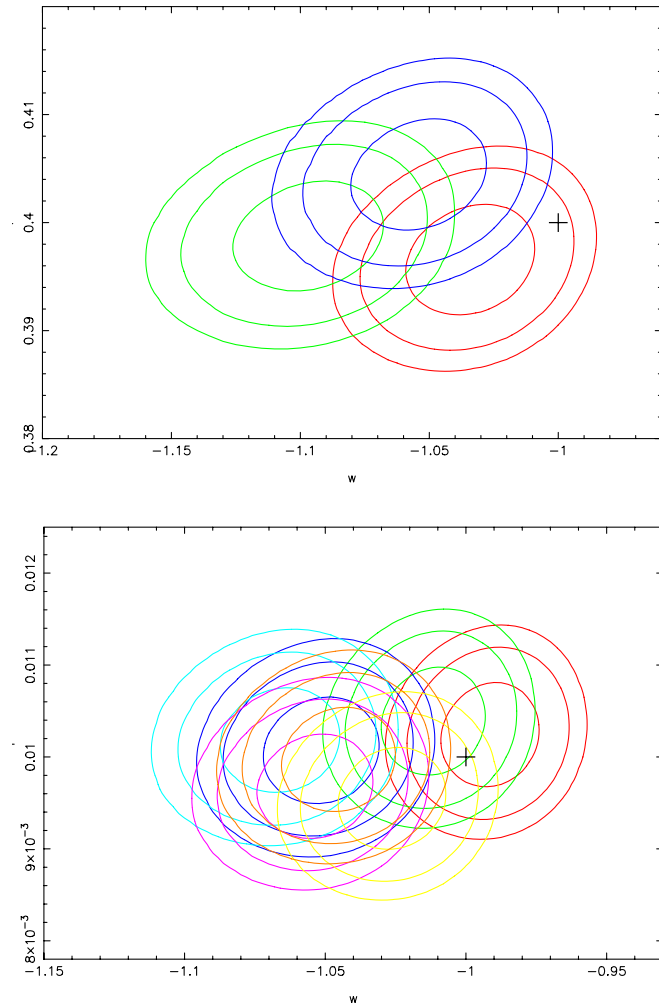


FIG. 4 (color online). Parameter constraints obtained from full-sky lensed simulations using the B -mode power spectrum only and assuming Gaussianity. Contours are shown at the 67%, 95% and 99% confidence limits. Top: Model-A simulations. Bottom: Model-B simulations. The parameters in the fiducial model are marked with a cross.

IV. NON-GAUSSIAN COVARIANCE

Equation (3) shows that the power in lensed B -modes on any particular scale arises from the power on a range of scales in E and ϕ . As can be seen in Fig. 1, the power in the lensing deflection field peaks at large scales, meaning that a significant fraction of the power in B -modes on scales around the peak is generated by large-scale lenses (see also Fig. 3). A particular mode in the lensing deflection field will lens many different E -modes to generate B -modes on a range of scales. The power in these B -modes will be related to the power in the single mode of the deflection field and so the measured power in B -modes is significantly correlated between scales [13]. In order to account for this non-Gaussianity in our analysis, we need to calculate the covariance between the measured \hat{C}_ℓ^{BB} . The covariance can be expressed as a sum of the Gaussian and non-Gaussian parts:

$$\text{Cov}(\hat{C}_{\ell_i}^{BB}, \hat{C}_{\ell_j}^{BB}) = S_{ij} = S_{ij}^G + S_{ij}^{\text{NG}}, \quad (9)$$

where the Gaussian part arises from the nonconnected part of the B -mode four-point function and on the full sky is given by

$$S_{ij}^G = \delta_{ij} \frac{2(C_{\ell_i}^{BB})^2}{2\ell + 1}. \quad (10)$$

Lensing has only a very small effect on any primordial B -modes from gravitational waves so we can treat these as an independent Gaussian field; their power spectrum then only enters the Gaussian part of the covariance matrix. Similar comments would apply to Gaussian instrument noise.

We can calculate the non-Gaussian part of the covariance matrix on the flat sky using the gradient approximation Eq. (2). The gradient approximation is not suitable for accurate calculation of the lensed power spectra [14]. In addition, the flat-sky approximation introduces percent level errors in the lensed B -mode spectrum on all scales [20]. Fortunately, the errors from these approximations tend to cancel above $\ell \sim 400$ and the lensed B -mode spectrum can be computed to $\sim 1\%$ accuracy on such scales using Eq. (3). We anticipate that the non-Gaussian components of the covariance matrix computed with the gradient and flat-sky approximations will have a similar level of accuracy.

We thus calculate the non-Gaussian part of the covariance matrix within the flat-sky approximation using the bandpower expression given in Ref. [13]:

$$\begin{aligned}
S_{ij}^{\text{NG}} = & \frac{2}{A\alpha_i\alpha_j} \int_{\ell \in i} d^2\ell_i \int_{\ell \in j} d^2\ell_j \int \frac{d^2\mathbf{L}}{(2\pi)^2} \{ W^2(\ell_i, \ell_i - \mathbf{L}) W^2(\ell_j, \ell_j - \mathbf{L}) C_{|\ell_i - \mathbf{L}|}^{\tilde{E}\tilde{E}} C_{|\ell_j - \mathbf{L}|}^{\tilde{E}\tilde{E}} (C_L^{\phi\phi})^2 \\
& + W^2(\ell_i, \mathbf{L}) W^2(\ell_j, \mathbf{L}) (C_L^{\tilde{E}\tilde{E}})^2 C_{|\ell_i - \mathbf{L}|}^{\phi\phi} C_{|\ell_j - \mathbf{L}|}^{\phi\phi} + W(\ell_i, \ell_i - \mathbf{L}) W(-\ell_i, \ell_j - \mathbf{L}) W(\ell_i, \ell_j - \mathbf{L}) W(-\ell_j, \ell_i - \mathbf{L}) \\
& \times C_{|\ell_i - \mathbf{L}|}^{\tilde{E}\tilde{E}} C_{|\ell_j - \mathbf{L}|}^{\tilde{E}\tilde{E}} C_L^{\phi\phi} C_{|\ell_i + \ell_j - \mathbf{L}|}^{\phi\phi} \}, \quad (11)
\end{aligned}$$

where A is the area of sky observed, $\alpha_i = \int_{\ell \in i} d^2\ell$ and similarly for α_j , and W is given by Eq. (4). This expression assumes periodic boundary conditions both to remove the complications due to E - B mixing and the additional correlations between the *observed* Fourier modes due to the geometry. For surveys of a few-hundred deg^2 , the former should be negligible except on the largest scales, and the latter can be dealt with by choosing wide enough bands. Neither of these is an issue for the full-sky observations we consider here, and in this case we can set $A = 4\pi$ and use bands of width $\Delta\ell = 1$. This means that we can simplify this integral somewhat, since for a fixed length of $\ell_i = \ell_i$ and $\ell_j = \ell_j$ the value of the integrand is dependent only on the angle between ℓ_i and ℓ_j and their moduli. Therefore we can take out the integral over ℓ_j and instead fix $\ell_j = (\ell_j, 0)$. We can also simplify the integral over ℓ_i to a one-dimensional integral in θ . We can summarize this as

$$\frac{2}{A\alpha_i\alpha_j} \int_{\ell \in i} d^2\ell_i \int_{\ell \in j} d^2\ell_j \rightarrow \frac{1}{4\pi^2} \int_0^{2\pi} d\theta_i. \quad (12)$$

Figure 5 illustrates the structure of the non-Gaussian component of the covariance matrix as calculated for




FIG. 5 (color online). Illustration of the non-Gaussian component of covariance matrix $\text{Cov}(\hat{C}_{\ell_1}^{BB}, \hat{C}_{\ell_2}^{BB})$. The maximum value at the peak is $5 \times 10^{-16} \mu\text{K}^4$ but the scale has been set to illustrate the structure in the rest of the matrix more clearly.

Model A. It can be seen that the correlation is present across all scales and is not confined to the region near the diagonal. Although the non-Gaussian covariance between any two \hat{C}_{ℓ}^{BB} is small compared to the Gaussian covariance of the diagonal elements ($\sim 3 \times 10^{-12} \mu\text{K}^4/\ell$ for $\ell \lesssim 300$), the fact that the non-Gaussian covariance spans a wide range of scales means that its net effect is very significant.

It is computationally expensive (taking several hours of CPU time) to recalculate the covariance matrix for each set of model parameters as would be required in parameter estimation. However, when considering variations in those parameters that are not well constrained by the primary anisotropies (w and $\Omega_\nu h^2$ here), the main effect on the lensed B -mode power spectrum is an overall scaling in amplitude. For this reason, we can approximate the covariance as

$$S_{ij} = S_{ij}^{\text{G}} + \frac{C_{\ell_i}}{C_{\ell_i}^{\text{fid}}} \frac{C_{\ell_j}}{C_{\ell_j}^{\text{fid}}} S_{ij}^{\text{NG, fid}}, \quad (13)$$

where C_{ℓ} is the lensing contribution to the B -mode power spectrum and C_{ℓ}^{fid} is the same in the fiducial model. This is similar to the scheme suggested in Ref. [13].

We can calculate the theoretically-optimum constraints on parameters θ_i by performing a Fisher analysis. For Gaussian fields, the Fisher information matrix reduces exactly to [30]

$$F_{ij} = \sum_{\ell=2}^{\ell_{\text{max}}} \sum_{X,Y} \frac{\partial C_{\ell}^X}{\partial \theta_i} \text{Cov}^{-1}(C_{\ell}^X, C_{\ell}^Y) \frac{\partial C_{\ell}^Y}{\partial \theta_j}, \quad (14)$$

where X and Y can be TT , EE , TE or BB and $\text{Cov}^{-1}(C_{\ell}^X, C_{\ell}^Y)$ is the inverse covariance matrix of the various measured power spectra. If we wish to consider the information that can be obtained from the lensed B -mode power spectrum, we need to compute the Fisher matrix $-\langle \partial^2 \ln P / \partial \theta_i \partial \theta_j \rangle$ where $P = P(\hat{C}_{\ell}^{BB} | \boldsymbol{\theta})$. If we can approximate this as Gaussian in the measured power spectrum, then provided the number of independent degrees of freedom is much larger than unity⁷ we find

⁷It is in just this limit of a large number of degrees of freedom that the likelihood will approach a Gaussian.

$$F_{ij} \approx \sum_{\ell\ell'} \frac{\partial C_{\ell}^{BB}}{\partial \theta_i} S_{\ell\ell'}^{-1} \frac{\partial C_{\ell'}^{BB}}{\partial \theta_j}, \quad (15)$$

where C_{ℓ}^{BB} is the total B -mode power spectrum.

The tightest parameter constraints achievable under any circumstances from CMB data alone are those obtained from the power spectra of the unlensed, Gaussian, temperature and polarization fields together with the lensing deflection field. Some approximation to this situation may be achievable using the full likelihood function for the lensed CMB fields [17,18], or, with larger errors, from quadratic reconstruction techniques [15,16]. In the top plot of Fig. 6 we show the constraints obtained from a Fisher analysis of Model A, if all parameter values are fixed except r and w . We consider three cases: using only the lensed B -mode spectrum with the correct covariance matrix; the same but (wrongly) using only the Gaussian part of the covariance matrix; and using all of the unlensed CMB fields and the lensing deflection field (treated as Gaussian). In all cases we only kept modes up to $\ell = 1200$. In the latter case, we include the unlensed CMB temperature and electric polarization though these have very little impact since the parameters we allow to vary are chosen to be nearly degenerate with respect to these fields. If the non-Gaussianity is neglected the constraints on w from the B -mode power spectrum alone are slightly tighter than those obtainable using all of the unlensed fields, which shows immediately that we are making a false assumption [12]. It can be seen that the errors on r are almost unaffected by the non-Gaussianity whereas the errors on w are significantly affected. However the presence of the lensed B -modes does significantly increase the errors on r [31,32], so if the gravitational-wave amplitude is small ($r \lesssim 0.01$) it will be desirable to attempt reconstruction of the unlensed CMB in order to detect the signal.

The degeneracy between w and $\Omega_{\nu} h^2$ that is present when only the B -mode power spectrum is used can be clearly seen in the bottom plot of Fig. 6. Here, r , w and $\Omega_{\nu} h^2$ are allowed to vary around a fiducial model with the same parameters as Model A except for a smaller $r = 0.005$ and massive neutrinos with $\Omega_{\nu} h^2 = 0.004$. It should be recognised that a Fisher analysis, based on the derivatives of the power spectra around the fiducial model values, is good at highlighting degeneracies but does not give accurate error contours for degenerate parameters [23]. The actual degeneracy between w and $\Omega_{\nu} h^2$ follows a curve rather than a straight line. The degeneracy is broken by using the power spectrum of the lensing deflection field, although even then the variables are still correlated to a lesser extent. In the unlensed, foreground-free ideal case, the errors on r are proportional to r and so there is no lower limit to the value of r that can be detected. In practice foreground residuals (and other systematic effects) are expected to pose a significant challenge and may well be




FIG. 6 (color online). 95% confidence limits from a Fisher analysis. Top: Model A. Bottom: constraints for a model with same parameters as Model A except $r = 0.005$ and $\Omega_{\nu} h^2 = 0.004$ (which we also allow to vary in the analysis). Blue dashed contours: limits if the lensed B -mode power spectrum is used, with the correct covariance matrix. Green solid line (top plot only): limits using the B -mode power spectrum and incorrectly assuming the data are Gaussian. Red filled area: theoretical limits if all unlensed power spectra are used including the power spectrum of the lensing deflection field.

the factor that limits the constraints that can ultimately be achieved.

V. NON-GAUSSIAN LIKELIHOOD

Having shown clearly that the non-Gaussianity of lensed B -modes cannot be ignored, we now address the question of how it may correctly be taken into account during parameter estimation from measured power spectra. To do this, we need to employ a likelihood function which includes information about the dependencies between the power on different scales. In principle, the exact likelihood

can be found from that for the lensed fields [17,18], but this would be very complicated, and such an analysis would probably be no simpler than an optimal analysis with the fields themselves (which would avoid the lossy compression to power spectra). Instead we proceed heuristically, modelling the likelihood as a function of the lensed spectrum and its covariance only. This necessarily assumes that the connected moments of the lensed field can be ignored above the four-point level, or, more correctly, that they can be approximated in a hierarchical manner by the lower moments. We only consider noise-free analysis on the full sky here, or, in Sec. VI, on a periodic flat sky. Since our approach is based on the measured power spectrum, the additional (geometric) complications from the survey geometry and those from instrument noise can be dealt with by standard techniques (see e.g. [33]).

A suitable likelihood function must reduce to a good approximation to the true likelihood in the limit that the non-Gaussian covariance tends to zero, Eq. (8). Note that it is the dependence of this function on the cosmological parameters θ that is of interest. The coupling of power between different scales by non-Gaussianity has a similar effect to the geometric coupling between scales that arises in observations of Gaussian fields over only part of the sky. A number of approximate likelihoods have been suggested for this latter problem, and a simple strategy for analyzing lensed B -mode spectra might be to replace the (geometric) covariance matrix in these approximations by the non-Gaussian covariance matrix.

A possible first approximation, a likelihood function that is Gaussian in \hat{C}_ℓ , is known to produce biased parameter constraints from Gaussian fields [34]. A better approximation is the log-normal distribution in which the likelihood is Gaussian in the log of the power. This was originally developed for the case where the peak of the likelihood (considered as a function of the theory C_ℓ ⁸) and its curvature there are known [34], but it can easily be tailored to our problem by replacing the modal C_ℓ by the measured \hat{C}_ℓ , and the (inverse) curvature by the theoretical covariance matrix. (This likelihood was denoted by \mathcal{L}'_{LN} in Ref. [35], with the prime distinguishing it from the original formulation in Ref. [34]; we make the same distinction in Eq. (16) below.) The log-normal likelihood has been employed in the analysis of several different CMB data sets, e.g. [36–40]. However, this likelihood is also biased; the level of bias is acceptable for many data sets but becomes significant when using full-sky data with a low noise level. The first-year WMAP data were analyzed using a likelihood function which is a weighted combination of the Gaussian and log-normal likelihoods [35]:

$$\ln P(\hat{C}|\theta) = \frac{1}{3} \ln P_{\text{Gauss}}(\hat{C}|\theta) + \frac{2}{3} \ln P'_{\text{LN}}(\hat{C}|\theta). \quad (16)$$

⁸For Gaussian fields, the parameters θ only enter the likelihood through the power spectrum.

This is a significantly better approximation to the exact likelihood in the limit of no non-Gaussianity; see Fig. 10 of the Appendix for comparisons.

To illustrate the level of bias of the different likelihood functions when estimating r and w , we created a Gaussian simulation of the CMB using the lensed B -mode power spectrum, and analyzed the resulting measured power spectrum. The results are shown in Fig. 7. It can be seen that the constraints obtained when the WMAP likelihood function is employed are virtually indistinguishable from those obtained using the exact likelihood given by Eq. (8), although the Gaussian and log-normal approximations do lead to significant differences.

Figure 8 shows the marginalized distribution of w obtained when the WMAP likelihood function is used to estimate parameters from the ten lensing simulations of Fig. 4. The non-Gaussian covariance matrix was included with the prescription in Eq. (13). Although from the spread of the maximum-likelihood values it seems that the width of the distributions are a good reflection of the uncertainties in the parameters, there is a significant bias towards low values of w , which means that, for most of the simulations, the parameter estimates are inconsistent with the input values. The marginalized distributions of r are very similar to those obtained when Eq. (8) was applied to the same simulations. This indicates that the non-Gaussian contribution to the covariance is less significant on large scales, and hence the presence of non-Gaussianity does not




FIG. 7 (color online). Likelihood contours at 67, 95 and 99% confidence limits from analysis of a Gaussian simulation of the CMB made using the lensed power spectrum. Red solid contours: using the exact likelihood function. Green dotted contours: using the Gaussian approximation to the likelihood function, showing the significant difference from the exact likelihood. Light blue dot-dashed contours: using the log-normal distribution, which gives a smaller, and opposite, shift. Dark-blue dashed contours: using the WMAP likelihood function, showing almost indistinguishable results from the exact likelihood. The fiducial model values are marked with a cross.

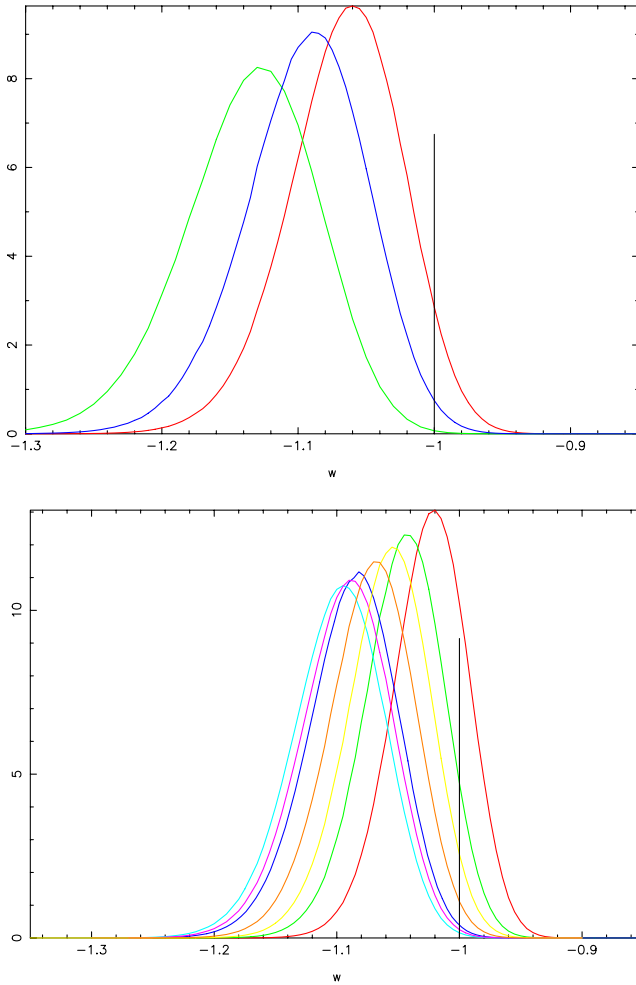


FIG. 8 (color online). Marginalized distributions of w obtained for three Model-A simulations (top) and seven Model-B simulations (bottom), using the WMAP likelihood function.

introduce a bias, or affect the errors, on r . This is understandable since the signal which contains the information about r is Gaussian to a high level.

VI. NEW LIKELIHOOD FUNCTION

In the previous section we showed that none of the existing likelihood approximations in the literature appear suitable for the analysis of future cosmic-variance limited, lensed B -mode data. For this reason we have re-examined the issue of likelihood approximations in CMB analysis; our findings can be found in the Appendix. The main result is a new likelihood function that appears to out-perform existing approximations. Here we summarize this new likelihood function and demonstrate that it produces accurate parameter constraints on simulated lensed data.

To begin, it is worth first reviewing the underlying issues that arise when constraining parameters from non-Gaussian fields. For ideal, full-sky observations, in the absence of non-Gaussianity, the probability of the observed

fields \mathbf{d} given the set of cosmological parameters $\boldsymbol{\theta}$ can be written as a function of the theoretical power spectrum $C_\ell(\boldsymbol{\theta})$ and the measured power spectrum \hat{C}_ℓ :

$$P(\mathbf{d}|\boldsymbol{\theta}) = P(\mathbf{d}|C_\ell) = f(\hat{C}_\ell)P(\hat{C}_\ell|C_\ell), \quad (17)$$

where $f(\hat{C}_\ell)$ is a function of the measured power spectrum only. It follows that \hat{C}_ℓ are sufficient statistics for parameter estimation and hence there is no loss of information in compressing the data to the measured power spectrum. For non-Gaussian fields, even for ideal observations, the measured power is no longer a sufficient statistic and the likelihood may not just be a function of the theoretical power spectra of the fields. Nevertheless, we can regard the measured power spectrum as a form of lossy compression of the data, and the relevant likelihood is then the sampling distribution $P(\hat{C}_\ell|\boldsymbol{\theta})$.

In the Appendix we derive a new likelihood by approximating the properly-normalized distribution $P(\hat{C}_\ell|\boldsymbol{\theta})$ as Gaussian in some function of the \hat{C}_ℓ . The likelihood takes the form

$$\ln P(\hat{C}_\ell|\boldsymbol{\theta}) \approx \ln A - \frac{1}{2} \sum_{\ell\ell'} M_{\ell\ell'}^{-1} (\hat{x}_\ell - \mu_\ell)(\hat{x}_{\ell'} - \mu_{\ell'}), \quad (18)$$

where

$$\hat{x}_\ell = \hat{C}_\ell^{1/3} \quad (19)$$

$$\mu_\ell = \left(\frac{2\ell - 1}{2\ell + 1} C_\ell \right)^{1/3}, \quad (20)$$

and

$$M_{\ell\ell'}^{-1} = 3C_\ell^{2/3} \left(\frac{2\ell - 1}{2\ell + 1} \right)^{1/6} S_{\ell\ell'}^{-1} 3C_{\ell'}^{2/3} \left(\frac{2\ell' - 1}{2\ell' + 1} \right)^{1/6}, \quad (21)$$

where $S_{\ell\ell'}$ is the covariance matrix of the measured \hat{C}_ℓ at parameters $\boldsymbol{\theta}$. The normalization is

$$A^{-1} \propto \sqrt{\det M_{\ell\ell'}} \prod_{\ell} \mu_\ell^2, \quad (22)$$

which we approximate here by $A \propto \prod_{\ell} 1/C_\ell$. Note that the normalization depends on the cosmological parameters and including this dependence is important to get an accurate approximation for the $\boldsymbol{\theta}$ -dependence of the likelihood in the tails.

In the Appendix we show that this new likelihood is considerably more accurate than other approximations in current use when tested on full-sky, Gaussian data when the number of degrees of freedom is low (i.e. low ℓ). When applied to Gaussian simulations, we find parameter constraints very similar to those shown for the WMAP (and exact) likelihood functions in Fig. 7. However, the new likelihood produces very different results when applied to non-Gaussian lensed simulations with the non-Gaussian covariance employed, as shown in Fig. 9. It can be seen here that for all three Model-A simulations the parameters




FIG. 9 (color online). A comparison of the parameter constraints obtained for the three Model-A simulations, using different likelihood functions, with contours shown at 67, 95 and 99% confidence limits. Red dot-dash contours: constraints obtained assuming Gaussianity. Green dashed contours: using WMAP likelihood function. Blue solid contours: using new likelihood function. The fiducial model value is shown with a cross.

constraints obtained are consistent with the values used in the simulations, unlike those obtained when the WMAP likelihood function is used. The results from the Model-B simulations are similar. In all cases the marginalized constraints on r are nearly identical to those obtained when Gaussianity is assumed, which indicates that the non-Gaussianity does not noticeably affect the estimated value of r , although the presence of the lensed power spectrum itself does influence the bounds we can place on r . The parameter constraints are consistent with the theoretical errors obtained from the Fisher analysis shown in Fig. 6.

In order to test the new likelihood function further, we performed flat-sky lensed simulations, which can be computed much more rapidly than the full-sky simulations. It was found that, with a pixel size of 1.3 arcmin, accurate power spectra could be obtained by interpolation provided that a cubic interpolation method was used (a simple linear interpolation resulted in an inaccurate lensed power spectrum). Periodic boundary conditions were specified. The amplitude of the lensed B -mode power spectrum does not vary linearly with w , and for measured power spectra with large uncertainties, this results in a distribution of w which is highly skewed and stretches far into the region $w < -1$. To avoid this, we used maps which were 91.7° across to reduce the sample variance. It should be noted that for this size of map approximating the sky as flat is clearly incorrect; the statistics of the lensed modes in the simulations will differ at the percent level from the spherical expectations on all scales [20], and more so for those modes approaching the survey size. We performed 150 of these simulations, with Model-A parameters.

The estimated power spectrum of the flat-sky maps is calculated as bandpowers:

$$\hat{C}_i = \frac{1}{A} \frac{\sum_{\ell \in i} \ell^2 |B(\ell)|^2}{2\pi \sum_{\ell \in i}}, \quad (23)$$

where $B(\ell)$ is the measured B -mode, the sum is over values of ℓ that lie in band i , and, recall, A is the area of the sky. In the limit that $\ell^2 C_\ell / (2\pi)$ is constant within each band, we find that $\langle \hat{C}_i \rangle = \ell^2 C_{\ell_i} / (2\pi)$. The WMAP likelihood function can straightforwardly be generalized to work with bandpowers. For the new likelihood function we need to drop the $(2\ell - 1)/(2\ell + 1)$ factors; they are more significant at low- ℓ where the flat-sky approximation does not hold anyway. The bandpower version of the new likelihood function can then be expressed as

$$P(\hat{C}|\theta) \propto \frac{1}{\prod_B C_B} \exp \left[-\frac{1}{2} \sum_{BB'} (\hat{C}_B^{1/3} - C_B^{1/3}) 3C_B^{2/3} \mathcal{S}_{BB'}^{-1} 3C_{B'}^{2/3} \times (\hat{C}_{B'}^{1/3} - C_{B'}^{1/3}) \right], \quad (24)$$

where the Gaussian part of the bandpower covariance matrix $\mathcal{S}_{B_i B_j}$ is given by

$$S_{B_i B_j}^G = \delta_{ij} \frac{2(2\pi)^2}{A\alpha_i^2} \int_{\ell \in i} d^2\ell \left(\frac{\ell^2}{2\pi} C_\ell^{BB} \right)^2, \quad (25)$$

and the non-Gaussian part can be calculated from the non-Gaussian part of the the full-sky covariance matrix as:

$$S_{B_i B_j}^{NG} = \frac{4\pi}{A} \frac{1}{\alpha_i \alpha_j} \sum_{\ell_i \in i} 2\pi\ell_i \sum_{\ell_j \in j} 2\pi\ell_j \frac{\ell_i^2 \ell_j^2}{(2\pi)^2} S_{\ell_i \ell_j}^{NG, \text{full}}. \quad (26)$$

Since the posterior distribution of w is skewed, as mentioned above, the maximum-likelihood value is greater than the mean value. Using the WMAP likelihood function, both the mean and maximum-likelihood values of w were significantly biased, with values (averaged over the simulations) of -1.075 ± 0.009 and -1.044 ± 0.007 respectively. The quoted errors are standard errors in the mean of 150 simulations, and were estimated from the simulations. For both the mean and maximum-likelihood value, the bias is comparable to the random error on the measured value; the situation would worsen if we considered a larger survey area as the random error would fall but the bias would remain. Using the new likelihood, the mean of w was -0.992 ± 0.007 and the maximum-likelihood value -1.017 ± 0.007 , showing that this likelihood function performs significantly better. A comparison between the width of the marginalized distributions of w and the spread of the mean values estimated from each of the flat-sky simulations gave good agreement, to around one percent. This shows that the estimated errors on the value of w are correct once the non-Gaussian covariance is included. There was a slight positive bias in the mean of r of about 0.5% for the new likelihood and around twice this for the WMAP likelihood. This is about 30% of the random error on r for the sample-variance limited observations considered here, but we expect that the bias is an artifact of the flat-sky simulations on large scales.

VII. CONCLUSIONS

The generation of B -mode polarization by weak gravitational lensing provides a way of measuring cosmological parameters from the CMB alone that would otherwise be poorly constrained. We showed that a power-spectrum analysis is wasteful, and, in particular, suffers from a strong degeneracy between the equation of state of the dark-energy and neutrino masses. A more optimal analysis that reconstructs the lensing deflection field can break this degeneracy [27], as can external data. Nevertheless, compression into the measured power spectrum will be a useful first step in the analysis of future B -mode data as it is likely to be more robust against real-world effects than more optimal reconstruction techniques. With this in mind, we have shown that it will be essential to take the non-Gaussianity of the lens-induced B -modes into account in a high signal-to-noise power-spectrum analysis, not only to avoid under-estimation of errors [13] but also to remove

biases in the parameters that affect lensing; constraints on the gravitational-wave amplitude are not noticeably affected by non-Gaussianity. Including the non-Gaussian covariance of the measured B -mode power spectrum in existing likelihood functions that have been used for parameter estimation from CMB spectra was shown to give biased results on noise-free simulations of lensed CMB fields. To remedy this we developed a new likelihood function that is more accurate in the tails of the distribution when the number of independent modes is low. Because of the non-Gaussian nature of CMB lensing, the number of such independent modes below a given multipole is lower than for Gaussian fields. We verified on simulations that the new likelihood function performs much better than existing approximations for the particular application considered here. However, we expect that it will be more generally applicable for the accurate analysis of the spectra of Gaussian CMB fields on large scales.

ACKNOWLEDGMENTS

The authors would like to thank Antony Lewis for many helpful comments throughout the course of this work. S S was supported by PPARC and the Astrophysics Group, Cavendish Laboratory and GR by the Leverhulme Trust for the duration of this work. AC is supported the Royal Society. Some of the results in this paper have been derived using the HEALPix package [43].

APPENDIX: LIKELIHOOD APPROXIMATIONS

In this appendix we derive the likelihood approximation employed in Sec. VI. Our strategy for dealing with the non-Gaussianity of lens-induced B -mode data in this paper is heuristic: we include the non-Gaussian correlation between power on different scales but our choice of likelihood is motivated by the analysis of Gaussian fields. For this reason, our new likelihood approximation should also be useful more generally in CMB analysis. We only consider observations of a single field here (\hat{C}_ℓ^B); we plan to address the more general problem of joint analysis of correlated fields in a future publication.

We develop two approximations motivated by the two ways that CMB power spectra are usually obtained. In the first, we work with the Gaussian CMB field \mathbf{d} directly, and aim to characterise the probability $P(\mathbf{d}|C_\ell)$ as a function of the theoretical power spectrum C_ℓ that encodes all of the cosmological information θ . Exploring $P(\mathbf{d}|C_\ell)$ is computationally expensive so often the modal value of C_ℓ and the curvature (or a Fisher approximation to it) are found, e.g. Ref. [41], and used in an analytic approximation to the C_ℓ -dependence of $P(\mathbf{d}|C_\ell)$. Complications due to, for example, the survey geometry, are accounted for only through their impact on the curvature matrix and modal C_ℓ . We construct an approximate likelihood by looking for variables $x_\ell(C_\ell)$ in which the exact likelihood is accurately

represented by a Gaussian, i.e. we approximate the C_ℓ -dependence of $\ln P(\mathbf{d}|C_\ell)$ as

$$\ln P(\mathbf{d}|C_\ell) \approx -\frac{1}{2} \sum_{\ell\ell'} M_{\ell\ell'}^{-1} (x_\ell - \mu_\ell)(x_{\ell'} - \mu_{\ell'}), \quad (\text{A1})$$

up to an irrelevant constant. If this approximation is to peak at the correct place, we require $\mu_\ell = x_\ell(C_{\ell,\text{ml}})$, where $C_{\ell,\text{ml}}$ is the modal (or maximum-likelihood) value of C_ℓ . Similarly, if the curvature at the peak is $-\mathcal{F}_{\ell\ell'}$, we must have

$$M_{\ell\ell'}^{-1} = \frac{dC_\ell}{dx_\ell} \frac{dC_{\ell'}}{dx_{\ell'}} \mathcal{F}_{\ell\ell'}, \quad (\text{A2})$$

where the derivatives are evaluated at the peak. To constrain the variable change $x_\ell(C_\ell)$, we examine the third derivatives of $\ln P(\mathbf{d}|C_\ell)$ with respect to the x_ℓ :

$$\begin{aligned} \frac{\partial^3 \ln P}{\partial x_\ell \partial x_{\ell'} \partial x_{\ell''}} &= \delta_{\ell\ell'} \delta_{\ell\ell''} \frac{\partial \ln P}{\partial C_\ell} \frac{d^3 C_\ell}{dx_\ell^3} \\ &+ \left(\delta_{\ell\ell'} \frac{\partial^2 \ln P}{\partial C_\ell \partial C_{\ell''}} \frac{d^2 C_\ell}{dx_\ell^2} \frac{dC_{\ell''}}{dx_{\ell''}} \right. \\ &\quad \left. + \text{cyclic perms.} \right) \\ &+ \frac{\partial^3 \ln P}{\partial C_\ell \partial C_{\ell'} \partial C_{\ell''}} \frac{dC_\ell}{dx_\ell} \frac{dC_{\ell'}}{dx_{\ell'}} \frac{dC_{\ell''}}{dx_{\ell''}}, \quad (\text{A3}) \end{aligned}$$

where we have imposed that x_ℓ is only a function of C_ℓ at the same ℓ . Setting the right-hand side exactly equal to zero is, in general, inconsistent. Instead, we obtain a tractable and consistent problem if we demand that it vanish everywhere when we replace the derivatives of P on the right with their expectation values evaluated with the likelihood appropriate to full-sky, noise-free observations [34]:

$$\ln P(\mathbf{d}|C_\ell) = -\sum_\ell (\ell + 1/2) (\hat{C}_\ell / C_\ell + \ln C_\ell), \quad (\text{A4})$$

where \hat{C}_ℓ is the measured C_ℓ . Evaluating the expectation values of the derivatives of this likelihood (over data from an ensemble with power C_ℓ), we find from Eq. (A3) that we must have

$$4(dC_\ell/dx_\ell)^2 = 3C_\ell d^2 C_\ell / dx_\ell^2. \quad (\text{A5})$$

Note that the full-sky number of degrees of freedom $2\ell + 1$ present in the derivatives cancels in this equation so we expect this procedure to remain valid for observations covering only a fraction f_{sky} of the sky. Equation (A5) is solved by $x_\ell = C_\ell^{-1/3}$, and gives an approximation for the C_ℓ -dependence of $\ln P(\mathbf{d}|C_\ell)$ of the form

$$\begin{aligned} \ln P(\mathbf{d}|C_\ell) &= -\frac{1}{2} \sum_{\ell\ell'} [3C_{\ell,\text{ml}}^{4/3} \mathcal{F}_{\ell\ell'} 3C_{\ell',\text{ml}}^{4/3} \\ &\quad \times (C_\ell^{-1/3} - C_{\ell,\text{ml}}^{-1/3})(C_{\ell'}^{-1/3} - C_{\ell',\text{ml}}^{-1/3})], \quad (\text{A6}) \end{aligned}$$

up to an irrelevant constant. The essential difference with

the log-normal approximation developed in Ref. [34] is that they demand that the expectation of the curvature of the likelihood be constant (when approximated by Eq. (A4)) whereas we force the expectation value of the derivative of the curvature to be zero. These are not equivalent since the latter preserves the distinction between $C_{\ell,\text{ml}}$ and C_ℓ until after the derivative of the curvature is taken while the former does not.

To include instrument noise, we repeat the above steps but now include isotropic noise in the full-sky likelihood in Eq. (A4). The effect in our approximate likelihood, Eq. (A6), is to add a noise offset N_ℓ to $C_{\ell,\text{ml}}$ and C_ℓ in the right-hand side. Following Ref. [34], we suggest determining N_ℓ by fitting the model

$$\mathcal{F}_{\ell\ell'}^{-1} \approx \frac{2}{(2\ell + 1)f_{\text{sky}}} (C_{\ell,\text{ml}} + N_\ell)^2, \quad (\text{A7})$$

to the diagonal elements of the inverse curvature at the peak.

The second approximation we consider is that which is used in Sec. VI. In this case, we first compress the data down to a measured power spectrum \hat{C}_ℓ and then ask what is the sampling distribution for the \hat{C}_ℓ given the cosmological model θ ? We then use the θ -dependence of this probability as the likelihood in parameter estimation [42]. For noise-free observations of Gaussian fields over the full sky, the sampling distribution is given by Eq. (8) which on taking logs becomes

$$\begin{aligned} \ln P(\hat{C}_\ell|C_\ell) &= -\sum_\ell \left(\frac{2\ell + 1}{2} \frac{\hat{C}_\ell}{C_\ell} + \frac{2\ell + 1}{2} \ln C_\ell \right. \\ &\quad \left. - \frac{2\ell - 1}{2} \ln \hat{C}_\ell \right), \quad (\text{A8}) \end{aligned}$$

up to a constant that is independent of \hat{C}_ℓ and the theoretical spectrum C_ℓ . This approach is more natural for the non-Gaussian problem considered in the body of this paper since there we perform (lossy) compression down to \hat{C}_ℓ^{BB} , and the power spectrum of the lensed B -modes does not fully characterize their statistics. We now look for a variable transformation $\hat{x}_\ell(\hat{C}_\ell)$ in which the sampling distribution $P(\hat{C}_\ell|\theta)$ is approximately Gaussian:

$$\ln P(\hat{C}_\ell|\theta) \approx \ln A - \frac{1}{2} \sum_{\ell\ell'} M_{\ell\ell'}^{-1} (\hat{x}_\ell - \mu_\ell)(\hat{x}_{\ell'} - \mu_{\ell'}), \quad (\text{A9})$$

where A is a \hat{C}_ℓ -independent (but cosmology-dependent) normalization. We assume that we can calculate the mean $\langle \hat{C}_\ell \rangle$ and covariance $S_{\ell\ell'}$ of the \hat{C}_ℓ given the cosmological model. Here we shall assume that \hat{C}_ℓ is constructed to be unbiased so $\langle \hat{C}_\ell \rangle = C_\ell$ although this assumption can easily be dropped. Since the mean and covariance of $P(\hat{C}_\ell|\theta)$ are specified rather than the modal value of \hat{C}_ℓ and the curvature there, we postpone their determination until we have determined the functional form of \hat{x}_ℓ . This is obtained by examining the third derivatives of the sampling distribution

with respect to \hat{C}_ℓ . These are given by Eq. (A3) but with \hat{C}_ℓ and \hat{x}_ℓ replacing C_ℓ and x_ℓ . Again we approximate the derivatives with respect to \hat{C}_ℓ , but this time with those derived from the distribution in Eq. (A8). The second and third derivatives are independent of C_ℓ and we make no further approximation with these. We ignore the first derivative term since it is small near the peak, so, on setting the (approximate) third derivative of $\ln P$ with respect to \hat{x}_ℓ to zero, we find

$$2(d\hat{C}_\ell/d\hat{x}_\ell)^2 = 3\hat{C}_\ell d^2\hat{C}_\ell/d\hat{x}_\ell^2. \quad (\text{A10})$$

This is solved by $\hat{x}_\ell = \hat{C}_\ell^{1/3}$. We could now determine μ_ℓ and $M_{\ell\ell'}$ by demanding that the approximate distribution gives the correct mean and covariance for \hat{C}_ℓ . A simpler method, which we adopt here, is to relate μ_ℓ and $M_{\ell\ell'}$ to the (known) mean and covariance using Eq. (A8) as a guide. For that ideal distribution, the peak is at $[(2\ell - 1)/(2\ell + 1)]\langle\hat{C}_\ell\rangle$ and the curvature at the peak is related to (minus) the inverse variance by a factor $[(2\ell + 1)/(2\ell - 1)]$. This motivates setting

$$\mu_\ell = \left(\frac{2\ell - 1}{2\ell + 1} C_\ell\right)^{1/3} \quad (\text{A11})$$

$$M_{\ell\ell'}^{-1} = \sqrt{\frac{2\ell + 1}{2\ell - 1} \frac{d\hat{C}_\ell}{d\hat{x}_\ell}} S_{\ell\ell'}^{-1} \sqrt{\frac{2\ell' + 1}{2\ell' - 1} \frac{d\hat{C}_{\ell'}}{d\hat{x}_{\ell'}}}, \quad (\text{A12})$$

where the derivatives are taken at the peak where $\hat{x}_\ell = \mu_\ell$. If we use this to approximate the exact result Eq. (A8) and evaluate the mean and standard deviation of \hat{C}_ℓ , we find

$$\frac{\langle\hat{C}_\ell\rangle}{C_\ell} = 1 + \frac{2}{27}\delta^2 + \dots \quad (\text{A13})$$

$$\sqrt{\frac{2\ell + 1}{2C_\ell^2}} \Delta\hat{C}_\ell = 1 + \frac{1}{18}\delta + \dots, \quad (\text{A14})$$

where $\delta \equiv (\ell - 1/2)^{-1}$. At $\ell = 2$ the fractional error on the mean and standard deviation are 3% and 4% respectively; by $\ell = 10$ these have dropped to 0.1% and 0.6%. It remains to fix the normalization. Exponentiating Eq. (A9) and integrating with the approximation $d\hat{C}_\ell \approx 3\mu_\ell^2 d\hat{x}_\ell$, we find that

$$A^{-1} \propto \sqrt{\det M_{\ell\ell'}} \prod_\ell \mu_\ell^2, \quad (\text{A15})$$

where the proportionality constant is independent of the cosmological parameters θ . In the text we further approximate the normalization by $A \propto 1/\prod_\ell C_\ell$.

To generalize our approximate sampling distribution to include instrument noise, we assume that the measured power spectrum has had a noise bias N_ℓ removed, so that $\langle\hat{C}_\ell\rangle = C_\ell$ still holds. The modifications we suggest are then $\hat{x}_\ell = (\hat{C}_\ell + N_\ell)^{1/3}$ and replacing C_ℓ by $C_\ell + N_\ell$ in μ_ℓ .

In Fig. 10 we compare our two new likelihood approximations with the exact distribution, and also with the log-normal distribution proposed in Ref. [34] and the likelihood used in the first-year WMAP analysis [35], for the case of noise-free, full-sky observations of Gaussian fields. In this case, the measured power spectrum \hat{C}_ℓ equals the




FIG. 10 (color online). Comparison of likelihood approximations for C_ℓ for noise-free, full-sky observations of Gaussian fields. Left: the exact distribution (red, solid), Eq. (A4); our new result (blue, dashed) in Eq. (A6) that assumes the peak position and curvature there; and our alternative result (green, dotted) in Eq. (A9) that works from the sampling distribution of the measured power spectrum, for $\ell = 2$ (left panel) and $\ell = 10$ (right panel). Note that the green dotted line is barely distinguishable from the exact distribution. Right: the log-normal distribution (blue, dashed); the WMAP likelihood (green, dotted) used in the analysis of the first-year data; and the exact likelihood (red, solid). In all cases $\epsilon \equiv C_\ell/\hat{C}_\ell$, where \hat{C}_ℓ is the measured C_ℓ ; \hat{C}_ℓ is also the maximum likelihood C_ℓ for the ideal conditions assumed.

maximum-likelihood spectrum $C_{\ell, \text{ml}}$, and working from the exact sampling distribution for \hat{C}_{ℓ} produces the same distribution for C_{ℓ} as working directly from the exact $P(\mathbf{d}|C_{\ell})$. Our new approximations are seen to be more accurate than the existing approximations, and particularly so when the number of degrees of freedom ($2\ell + 1$ in this example) is small. Our second approximation, based on the

sampling distribution of the measured C_{ℓ} , is seen to be more accurate in the tail of the distribution than our first approximation that starts from the peak position and curvature. This appears to be because of the nonperturbative way in which the theoretical C_{ℓ} appear in the normalization of the sampling distribution.

-
- [1] S. Church, P. Ade, J. Bock, M. Bowden, J. Carlstrom, K. Ganga, W. Gear, J. Hinderks, W. Hu, B. Keating *et al.* *New Astron. Rev.* **47**, 1083 (2003).
- [2] B. G. Keating, P. A. R. Ade, J. J. Bock, E. Hivon, W. L. Holzapfel, A. E. Lange, H. Nguyen, and K. W. Yoon, in *Proceedings of the SPIE, Polarimetry in Astronomy, 2003*, edited by S. Fineschi, vol. 4843, p. 284.
- [3] P. Oxley, P. A. Ade, C. Baccigalupi, P. deBernardis, H.-M. Cho, M. J. Devlin, S. Hanany, B. R. Johnson, T. Jones, A. T. Lee, *et al.*, in *Proceedings of the SPIE, Earth Observing Systems IX, 2004*, edited by and (2004), vol. 5543, p. 320.
- [4] A. C. Taylor, A. Challinor, D. Goldie, K. Grainge, M. E. Jones, A. N. Lasenby, S. Withington, G. Yassin, W. K. Gear, L. Piccirillo *et al.*, astro-ph/0407148.
- [5] M. Zaldarriaga, D. N. Spergel, and U. Seljak, *Astrophys. J.* **488**, 1 (1997).
- [6] M. Kamionkowski, A. Kosowsky, and A. Stebbins, *Phys. Rev. Lett.* **78**, 2058 (1997).
- [7] U. Seljak and M. Zaldarriaga, *Phys. Rev. Lett.* **78**, 2054 (1997).
- [8] M. Zaldarriaga and U. Seljak, *Phys. Rev. D* **58**, 023003 (1998).
- [9] U. Seljak, *Astrophys. J.* **463**, 1 (1996).
- [10] R. B. Metcalf and J. Silk, *Astrophys. J.* **489**, 1 (1997).
- [11] R. Stompor and G. Efstathiou, *Mon. Not. R. Astron. Soc.* **302**, 735 (1999).
- [12] W. Hu, *Phys. Rev. D* **65**, 023003 (2002).
- [13] K. M. Smith, W. Hu, and M. Kaplinghat, *Phys. Rev. D* **70**, 043002 (2004).
- [14] A. Lewis, *Phys. Rev. D*, **71**, 083008 (2005).
- [15] W. Hu, *Astrophys. J. Lett.* **557**, L79 (2001).
- [16] W. Hu and T. Okamoto, *Astrophys. J.* **574**, 566 (2002).
- [17] C. M. Hirata and U. Seljak, *Phys. Rev. D* **67**, 043001 (2003).
- [18] C. M. Hirata and U. Seljak, *Phys. Rev. D* **68**, 083002 (2003).
- [19] A. Challinor and G. Chon, *Phys. Rev. D* **66**, 127301 (2002).
- [20] A. Challinor and A. Lewis, *Phys. Rev. D* **71**, 103010 (2005).
- [21] G. Efstathiou and J. R. Bond, *Mon. Not. R. Astron. Soc.* **304**, 75 (1999).
- [22] U. Seljak, A. Makarov, P. McDonald, S. F. Anderson, N. A. Bahcall, J. Brinkmann, S. Burles, R. Cen, M. Doi, J. E. Gunn *et al.*, *Phys. Rev. D* **71**, 103515 (2005).
- [23] D. J. Eisenstein, W. Hu, and M. Tegmark, *Astrophys. J.* **518**, 2 (1999).
- [24] D. N. Spergel, L. Verde, H. V. Peiris, E. Komatsu, M. R. Nolta, C. L. Bennett, M. Halpern, G. Hinshaw, N. Jarosik, A. Kogut *et al.*, *Astrophys. J. Suppl. Ser.* **148**, 175 (2003).
- [25] Ø. Elgarøy and O. Lahav, *New J. Phys.* **7**, 61 (2005).
- [26] R. E. Smith, J. A. Peacock, A. Jenkins, S. D. M. White, C. S. Frenk, F. R. Pearce, P. A. Thomas, G. Efstathiou, and H. M. P. Couchman, *Mon. Not. R. Astron. Soc.* **341**, 1311 (2003).
- [27] M. Kaplinghat, L. Knox, and Y.-S. Song, *Phys. Rev. Lett.* **91**, 241301 (2003).
- [28] M. Zaldarriaga, *Phys. Rev. D* **62**, 063510 (2000).
- [29] J. E. Lidsey, A. R. Liddle, E. W. Kolb, E. J. Copeland, T. Barreiro, and M. Abney, *Rev. Mod. Phys.* **69**, 373 (1997).
- [30] M. Zaldarriaga and U. Seljak, *Phys. Rev. D* **55**, 1830 (1997).
- [31] L. Knox and Y.-S. Song *Phys. Rev. Lett.* **89**, 011303 (2002).
- [32] M. Kesden, A. Cooray, and M. Kamionkowski *Phys. Rev. Lett.* **89**, 011304 (2002).
- [33] A. Challinor, astro-ph/0502093.
- [34] J. R. Bond, A. H. Jaffe, and L. Knox *Astrophys. J.* **533**, 19 (2000).
- [35] L. Verde, H. V. Peiris, D. N. Spergel, M. R. Nolta, C. L. Bennett, M. Halpern, G. Hinshaw, N. Jarosik, A. Kogut, M. Limon *et al.* *Astrophys. J. Suppl. Ser.* **148**, 195 (2003).
- [36] J. H. Goldstein, P. A. R. Ade, J. J. Bock, J. R. Bond, C. Cantalupo, C. R. Contaldi, M. D. Daub, W. L. Holzapfel, C. Kuo, A. E. Lange *et al.* *Astrophys. J.* **599**, 773 (2003).
- [37] C. J. MacTavish, P. A. R. Ade, J. J. Bock *et al.*, astro-ph/0507503.
- [38] C. Pryke, N. W. Halverson, E. M. Leitch, J. Kovac, J. E. Carlstrom, W. L. Holzapfel, and M. Dragovan, *Astrophys. J.* **568**, 46 (2002).
- [39] A. Balbi, P. Ade, J. Bock *et al.*, *Astrophys. J. Lett.* **545**, L1 (2000).
- [40] J. A. Rubiño-Martín, R. Rebolo, P. Carreira *et al.*, *Mon. Not. R. Astron. Soc.* **341**, 1084 (2003).
- [41] J. R. Bond, A. H. Jaffe, and L. Knox, *Phys. Rev. D* **57**, 2117 (1998).
- [42] B. D. Wandelt, E. Hivon, and K. M. Górski, *Phys. Rev. D* **64**, 083003 (2001).
- [43] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann, *Astrophys. J.* **622**, 759 (2005).