Weak lensing of the CMB: Sampling errors on *B* modes

Kendrick M. Smith,[1] Wayne Hu,[2] and Manoj Kaplinghat[3]
[1]*Department of Physics, University of Chicago, Chicago, Illinois 60637, USA*
[2]*Center for Cosmological Physics, Department of Astronomy and Astrophysics, and Enrico Fermi Institute, University of Chicago, Chicago, Illinois 60637, USA*
[3]*Department of Physics, One Shields Avenue, University of California, Davis, California 95616, USA*
(Received 20 February 2004; published 9 August 2004)

The *B* modes generated by the lensing of cosmic microwave background (CMB) polarization are a primary target for the upcoming generation of experiments and can potentially constrain quantities such as the neutrino mass and dark energy equation of state. The net sample variance on the small scale *B* modes out to $l=2000$ exceeds Gaussian expectations by a factor of 10 reflecting the variance of the larger scale lenses that generate them. It manifests itself as highly correlated band powers with correlation coefficients approaching 70% for wide bands of $\Delta l/l \sim 0.25$. It will double the total variance for experiments that achieve a sensitivity of approximately 4 μK arcmin and a beam of several arcminutes or better. This non-Gaussianity must be taken into account in the analysis of experiments that go beyond first detection.

DOI: 10.1103/PhysRevD.70.043002 PACS number(s): 98.70.Vc, 98.80.Es

I. INTRODUCTION

As a step on the road toward the ultimate goal of detecting primordial gravitational waves, upcoming cosmic microwave background (CMB) polarization experiments will target the distortion to the acoustic polarization induced by gravitational lensing. As with the polarization induced by gravitational waves, the gravitationally lensed polarization contains a component with handedness, the so-called *B* mode component [1]. Unlike gravitational waves, gravitational lensing provides a guaranteed signal. In the standard cosmological model, the predicted amplitude of the *B* modes can only vary at the tens of percent level within current constraints [2]. Moreover these fine variations provide an opportunity to measure the dark side of the universe, namely the dark energy and neutrino dependent growth of structure, as well as another handle on the reionization optical depth [3,4].

Although both the intrinsic distribution and the density perturbations that lens the CMB are expected to be Gaussian, the lensed distribution is non-Gaussian at second order in the perturbations. The non-Gaussianity is therefore relatively small in the temperature distribution [5]. However because the *B* modes are generated by the lensing itself, its non-Gaussianity is a first order effect but fortunately one that is precisely calculable. Gravitational lensing therefore also provides a unique testing ground for experimentally extracting a non-Gaussian signal in the presence of foregrounds and systematic errors. Ultimately, the non-Gaussianity of the lensed polarization also provides the key to mapping the dark matter [6,7] and hence the separation of the lensing and gravitational wave *B* mode components [8,9].

For the upcoming generation of experiments, the non-Gaussianity will provide an important source of uncertainty for power spectrum measurements. Fisher information studies have shown that the information contained on cosmological parameters in the *B* mode power spectrum under the Gaussian approximation unphysically exceeds that contained in the two underlying Gaussian fields [3].

In this paper, we study the origin and quantify the impact of the non-Gaussian sample variance on *B* mode power spectra measurements. The basic reason for the large sample variance is that the fluctuations that lens the CMB are mainly on degree scales. All of the arcminute scale *B* modes fluctuate jointly with the lens and so precision measurements will require many degree scale patches not simply many arcminute scale patches.

We begin in Sec. II by briefly reviewing the generation of *B* modes through gravitational lensing. We calculate the non-Gaussian sample covariance in Sec. III and explore its impact on measurements in Sec. IV. We conclude in Sec. V. For illustrative purposes we employ throughout a fiducial cosmology that is consistent with Wilkinson Microwave Anisotropy Probe (WMAP) determinations: an initial scale invariant spectrum of curvature fluctuations with amplitude $\delta_\zeta = 5.07\times 10^{-5}$ ($\sigma_8=0.91$, $\tau=0.17$), a baryon density $\Omega_b h^2 = 0.024$ and a matter density $\Omega_m h^2 = 0.14$ in a flat $\Omega_\Lambda = 0.73$ cosmology.

II. *B* MODES

Weak lensing by the large-scale structure of the universe remaps the polarization field or equivalently the dimensionless Stokes parameters $Q(\hat{\mathbf{n}})$ and $U(\hat{\mathbf{n}})$ as [1,10,11]

$$[Q \pm iU](\hat{\mathbf{n}}) = [\tilde{Q} \pm i\tilde{U}][\hat{\mathbf{n}} + \nabla\phi(\hat{\mathbf{n}})], \tag{1}$$

where $\hat{\mathbf{n}}$ is the direction on the sky, tildes denote the unlensed field, and the deflection angle $\nabla\phi$ is the gradient of the line of sight projection of the gravitational potential $\Psi(\mathbf{x},D)$,

$$\phi(\hat{\mathbf{n}}) = -2\int dD \frac{(D_s - D)}{D D_s}\Psi(D\hat{\mathbf{n}}, D), \tag{2}$$

where D is the comoving distance along the line of sight in the assumed flat cosmology and D_s denotes the distance to the last-scattering surface. In the flat sky approximation, the Stokes parameters can be decomposed into *E* and *B* Fourier modes as

$$[Q \pm iU](\hat{\mathbf{n}}) = -\int \frac{d^2l}{(2\pi)^2}[E(\mathbf{l}) \pm iB(\mathbf{l})]e^{\pm 2i\varphi_l}e^{i\mathbf{l}\cdot\hat{\mathbf{n}}}, \quad (3)$$

where $\mathbf{l}=(l\cos\varphi_l, l\sin\varphi_l)$ and likewise

$$\phi(\hat{\mathbf{n}}) = \int \frac{d^2l}{(2\pi)^2}\phi(\mathbf{l})e^{i\mathbf{l}\cdot\hat{\mathbf{n}}}. \quad (4)$$

Even in the absence of an intrinsic $\tilde{B}$, lensing will generate B as [1]

$$B(\mathbf{l}) = \int \frac{d^2l'}{(2\pi)^2}\tilde{E}(\mathbf{l}')\phi(\mathbf{l}-\mathbf{l}')W(\mathbf{l},\mathbf{l}'), \quad (5)$$

where we have expanded Eq. (1) in the leading order gradient approximation and the mode coupling weight is

$$W(\mathbf{l},\mathbf{l}') = [\mathbf{l}'\cdot(\mathbf{l}-\mathbf{l}')]\sin 2(\varphi_l - \varphi_{l'}). \quad (6)$$

Upcoming experiments will mainly measure the power spectrum of the B modes. In the flat sky approximation, the power spectra of a statistically homogeneous field $X(\mathbf{l})$ is given by

$$\langle X^*(\mathbf{l})X(\mathbf{l}')\rangle = (2\pi)^2\delta(\mathbf{l}-\mathbf{l}')C_{\mathbf{l}}, \quad (7)$$

and statistical isotropy requires $C_{\mathbf{l}}=C_l$. It follows that (e.g., Ref. [12])

$$C_l^{BB} = \int \frac{d^2L}{(2\pi)^2}W^2(\mathbf{l},\mathbf{l}-\mathbf{L})C_{\mathbf{l}-\mathbf{L}}^{\tilde{E}\tilde{E}}C_L^{\phi\phi}, \quad (8)$$

given that the intrinsic polarization from $z\sim 1000$ is essentially uncorrelated with the lensing potential from $z\lesssim 3$. The B power spectrum is essentially a convolution of the $\tilde{E}$ and ϕ power spectra.

The total power in the B field arises from power in the ϕ field as

$$\Delta^2_{\rm tot} = \int \frac{d^2l}{(2\pi)^2}C_l^{BB} = (0.46\ \ \mu{\rm K}/T)^2$$
$$= \int \frac{d^2L}{(2\pi)^2}C_L^{\phi\phi}\left[\int \frac{d^2l}{(2\pi)^2}W^2(\mathbf{l},\mathbf{l}-\mathbf{L})C_{\mathbf{l}-\mathbf{L}}^{\tilde{E}\tilde{E}}\right] \quad (9)$$

where the value in parentheses is the rms in the fiducial cosmology and $T=2.725\times 10^6\ \mu$K is the CMB temperature. For reference, the rms of the $l\lesssim 2000$ low-pass filtered B field is 0.43 μK. We will typically take this value as the maximum l estimated for illustration purposes.

In Fig. 1, we show the B power spectrum in the fiducial cosmology. Note that the power spectrum peaks at $l\sim 10^3$ reflecting the power in the underlying unlensed $\tilde{E}$ power spectrum. However this mode coupling in Eq. (8) is achieved through power in the potential $C_L^{\phi\phi}$ across a broad range in L. The B field acquires half its power for $L\lesssim 450$. In other words, degree-scale gravitational lenses give rise to the B power at the $10'$ scale.

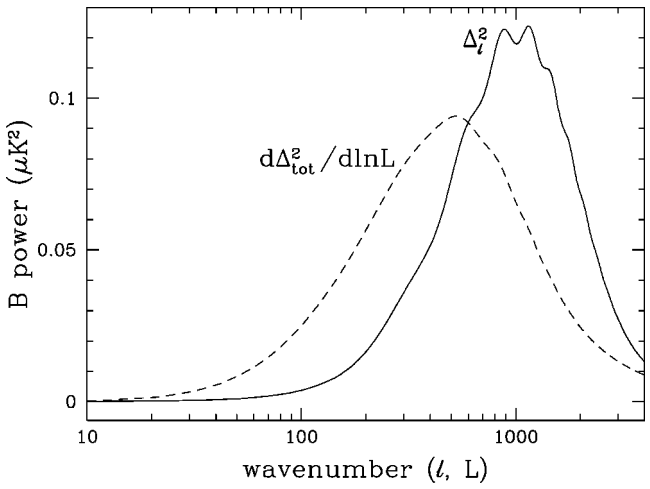

FIG. 1. B power spectrum $\Delta_l^2 = l^2C_l^{BB}/2\pi$ and the differential contribution to the total B variance $\Delta^2_{\rm tot}$ from power in the lensing field ϕ with wave number L [see Eq. (9)]. Even though the B power peaks at $l\sim 10^3$, it gains half its contribution from power in the lensing field at $L\lesssim 450$. Sampling errors in B therefore reflect the larger scale variations in ϕ.

This fact is the key to understanding the non-Gaussian covariance. In a given degree scale patch of sky there is effectively only a single lens contributing to the small scale B power. The amplitude of all of these smaller scale B modes will covary with the amplitude of the larger scale lens. Hence reducing the sampling errors for arcminute scale B modes requires many independent degree scale patches of sky. If the B field were Gaussian, such a constraint would require only an equivalent number of arcminute scale patches of sky.

We illustrate this effect with Monte Carlo realizations of the polarization field and lenses from the power spectra of the fiducial cosmology. The Monte Carlo simulations were performed using a square patch of sky of side length 22.9°, and with fields sampled on a grid with 1.3-arcmin spacing. Periodic boundary conditions were used to eliminate the boundary effects that would otherwise make the decomposition of Q and U into E and B ambiguous.

In Fig. 2, we show the band powers extracted from individual runs, using 13 bins logarithmically spaced from $l=100$ to $l=2000$. If the covariance were Gaussian, then each estimated band powers would be uncorrelated from bin to bin, with rms deviation given by the shaded band. The non-Gaussian covariance manifests itself here as a positive correlation from bin to bin, and a higher variance than one would expect from Gaussian statistics alone. This behavior is the Fourier analog of having all of the arcminute scale B modes fluctuate jointly in amplitude with the degree scale lenses.

III. SAMPLE COVARIANCE

On the large scales that are responsible for lensing, the potential fluctuations are nearly Gaussian, as are the unlensed E modes by assumption. The B-mode sample covariance can therefore be accurately quantified analytically.

Consider an ideal, noise-free estimator of the bandpower,

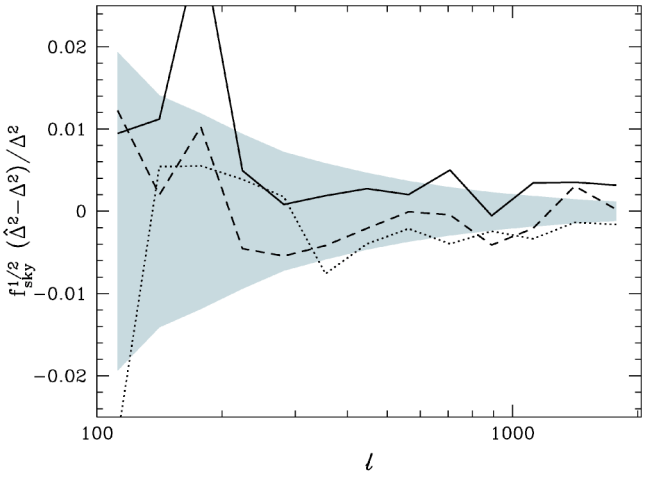

FIG. 2. Fractional deviation of band powers from the ensemble average for three Monte Carlo realizations (lines) compared with the expected rms deviation (shaded region) for a Gaussian field of the same power spectrum. The realizations show correlated deviations in the recovered band powers at high l.

$$\hat{\Delta}_i^2=\frac{1}{A\alpha_i}\int_{l\in i}d^2l\frac{l^2}{2\pi}B(\mathbf{l})B^*(\mathbf{l}), \tag{10}$$

where A is the survey area in steradians and

$$\alpha_i=\int_{l\in i}d^2l \tag{11}$$

is the l-space area of the band. For a flat spectrum

$$\Delta_l^2\equiv\frac{l^2C_l^{BB}}{2\pi}=\text{const}, \tag{12}$$

it is an unbiased estimator of the amplitude

$$\Delta_i^2\equiv\langle\hat{\Delta}_i^2\rangle=\Delta_l^2\,, \tag{13}$$

where we have used the relationship $(2\pi)^2\delta(\mathbf{0})=A$ for a finite patch of sky. In the limit of bandwidths approaching $\Delta l=1$, the relationship holds for any underlying power spectrum. The band power weight of $l^2/2\pi$ in Eq. (10) is appropriate near the peak at $l\sim10^3$. For the $l\ll10^3$ regime where the band power is steeply rising but the power is nearly flat (see Fig. 1), one can drop the weight factors here and still use wide bands.

Under the assumption of nonoverlapping bands in l, the sample covariance of the estimator then follows as

$$S_{ij}\equiv\langle(\hat{\Delta}_i^2-\Delta_i^2)(\hat{\Delta}_j^2-\Delta_j^2)\rangle=S_{ij}^G+S_{ij}^N. \tag{14}$$

The first piece is the Gaussian, or more properly the unconnected, contribution

$$S_{ij}^G=\delta_{ij}\frac{2(2\pi)^2}{A\alpha_i^2}\int d^2l_i\left(\frac{l_i^2}{2\pi}C_{l_i}^{BB}\right)^2. \tag{15}$$

In the limit of narrow bands and high $l\gg1$, the Gaussian piece takes a form that is familiar from Fisher matrix studies [13]

$$S_{ij}^G\approx\delta_{ij}\frac{2(2\pi)^2}{A\alpha_i}(\Delta_i^2)^2\approx\delta_{ij}\frac{2}{(2l_i+1)\Delta l_i f_{\rm sky}}(\Delta_i^2)^2, \tag{16}$$

where $f_{\rm sky}=A/4\pi$ is the fraction of sky covered. The Gaussian errors mainly reflect a mode counting argument. Since the fundamental mode $2\pi/A^{-1/2}$ sets the spacing of modes in l space, $A\alpha_j/(2\pi)^2$ is the total number of modes in the survey area [14].

The non-Gaussian or connected piece increases the band variances and correlates them

$$S_{ij}^N=\frac{2}{A\alpha_i\alpha_j}\int_{l_i\in i}d^2l_i\int_{l_j\in j}d^2l_j\int\frac{d^2L}{(2\pi)^2}\frac{l_i^2l_j^2}{(2\pi)^2}\times(a^{\mathbf{L}}_{\mathbf{l}_i\mathbf{l}_j}+b^{\mathbf{L}}_{\mathbf{l}_i\mathbf{l}_j}+c^{\mathbf{L}}_{\mathbf{l}_i\mathbf{l}_j})$$

where

$$\begin{aligned}
a^{\mathbf{L}}_{\mathbf{l}_i\mathbf{l}_j}&=W^2(\mathbf{l}_i,\mathbf{l}_i-\mathbf{L})W^2(\mathbf{l}_j,\mathbf{l}_j-\mathbf{L})C^{\tilde{E}\tilde{E}}_{\mathbf{l}_i-\mathbf{L}}C^{\tilde{E}\tilde{E}}_{\mathbf{l}_j-\mathbf{L}}(C_L^{\phi\phi})^2,\\
b^{\mathbf{L}}_{\mathbf{l}_i\mathbf{l}_j}&=W^2(\mathbf{l}_i,\mathbf{L})W^2(\mathbf{l}_j,\mathbf{L})(C_L^{\tilde{E}\tilde{E}})^2C^{\phi\phi}_{\mathbf{l}_i-\mathbf{L}}C^{\phi\phi}_{\mathbf{l}_j-\mathbf{L}},\\
c^{\mathbf{L}}_{\mathbf{l}_i\mathbf{l}_j}&=W(\mathbf{l}_i,\mathbf{l}_i-\mathbf{L})W(-\mathbf{l}_i,\mathbf{l}_j-\mathbf{L})W(\mathbf{l}_j,\mathbf{l}_j-\mathbf{L})\\
&\quad\times W(-\mathbf{l}_j,\mathbf{l}_i-\mathbf{L})C^{\tilde{E}\tilde{E}}_{\mathbf{l}_i-\mathbf{L}}C^{\tilde{E}\tilde{E}}_{\mathbf{l}_j-\mathbf{L}}C^{\phi\phi}_{\mathbf{L}}C^{\phi\phi}_{\mathbf{l}_i+\mathbf{l}_j-\mathbf{L}}.
\end{aligned} \tag{17}$$

Note that the both the Gaussian and non-Gaussian pieces are the same order in the $C_l^{\tilde{E}\tilde{E}}$ and $C_l^{\phi\phi}$ power spectra.

The non-Gaussianity is conveniently quantified by the variance degradation factor

$$D_i=\frac{S_{ii}}{S_{ii}^G}, \tag{18}$$

which gives the diagonals of the covariance matrix, and the correlation matrix

$$R_{ij}\equiv\frac{S_{ij}}{\sqrt{S_{ii}S_{jj}}}. \tag{19}$$

In Table I we show these quantities as calculated from the analytic expression (upper triangle) and 10^5 Monte Carlo simulations (lower triangle). With 10^5 iterations, the elements of R_{ij} have converged to a level ranging from around 0.01 in the lower bins, to around 0.001 in the higher bins (owing to the larger number of modes). The remaining discrepancies between the analytic and Monte Carlo results are at the 0.01 level and mainly reflect the use of the gradient approximation in Eq. (5) when deriving the analytic results; the Monte Carlo simulations resample Q and U as in Eq. (1). The close agreement between the two provides confirmation that the gradient approximation is accurate when computing covariances in the range $100\leqslant l\leqslant2000$. For these wide

TABLE I. Correlation matrix R_{ij} and variance degradation factor D_i. Parenthetical values are computed from 10^5 Monte Carlo simulations and compare well with the analytic calculation. The bands are chosen to be logarithmically spaced and nonoverlapping with $l_{\min i} = 0.89 l_i$ and $l_{\max i} = l_{\min i-1}$.

l_i	112	141	178	224	281	354	446	561	707	889	1120	1409	1775	D_i
112	1.00	0.03	0.04	0.05	0.05	0.06	0.06	0.06	0.07	0.07	0.07	0.08	0.08	1.03 (1.04)
141	(0.03)	1.00	0.05	0.06	0.07	0.07	0.08	0.08	0.08	0.09	0.09	0.10	0.10	1.04 (1.06)
178	(0.04)	(0.05)	1.00	0.09	0.10	0.09	0.10	0.11	0.11	0.12	0.13	0.13	0.13	1.07 (1.10)
224	(0.05)	(0.07)	(0.08)	1.00	0.13	0.13	0.14	0.15	0.15	0.17	0.17	0.18	0.18	1.13 (1.15)
281	(0.06)	(0.08)	(0.10)	(0.13)	1.00	0.18	0.18	0.20	0.21	0.23	0.23	0.24	0.24	1.21 (1.23)
354	(0.05)	(0.08)	(0.09)	(0.13)	(0.19)	1.00	0.22	0.24	0.27	0.27	0.28	0.29	0.29	1.28 (1.29)
446	(0.06)	(0.08)	(0.10)	(0.14)	(0.19)	(0.22)	1.00	0.29	0.30	0.32	0.32	0.33	0.32	1.33 (1.35)
561	(0.06)	(0.09)	(0.11)	(0.15)	(0.21)	(0.24)	(0.29)	1.00	0.37	0.40	0.40	0.40	0.39	1.55 (1.57)
707	(0.07)	(0.09)	(0.11)	(0.15)	(0.22)	(0.27)	(0.30)	(0.37)	1.00	0.47	0.48	0.48	0.47	1.76 (1.78)
889	(0.07)	(0.09)	(0.13)	(0.17)	(0.23)	(0.28)	(0.32)	(0.40)	(0.47)	1.00	0.56	0.56	0.55	2.17 (2.17)
1120	(0.07)	(0.10)	(0.13)	(0.17)	(0.24)	(0.28)	(0.32)	(0.40)	(0.47)	(0.55)	1.00	0.62	0.61	2.56 (2.55)
1409	(0.08)	(0.10)	(0.13)	(0.17)	(0.24)	(0.29)	(0.32)	(0.40)	(0.48)	(0.56)	(0.62)	1.00	0.66	2.98 (2.94)
1775	(0.08)	(0.10)	(0.13)	(0.17)	(0.23)	(0.29)	(0.32)	(0.40)	(0.46)	(0.54)	(0.60)	(0.66)	1.00	3.19 (3.17)

bands of $\Delta l/l \sim 0.25$, the non-Gaussian contribution triples the variance near the peak of the B power and correlates neighboring bands by 70%.

As the bands widen, the non-Gaussian effects superficially appear larger in the covariance matrix. In Fig. 3, we show the degradation factor D_i as a function of the bandwidth. The nearly linear scaling at high D_i can be understood as reflecting the linear decrease in the Gaussian errors with bandwidth in Eq. (16) revealing a non-Gaussian floor to the variance. For narrow bands, this effect is hidden in the band correlations. The net effect on the measurements is the same however: when combining narrow band measurements to constrain the parameters underlying the B power spectrum, their nonindependence leads to a large degradation in the constraints as we show in the next section.

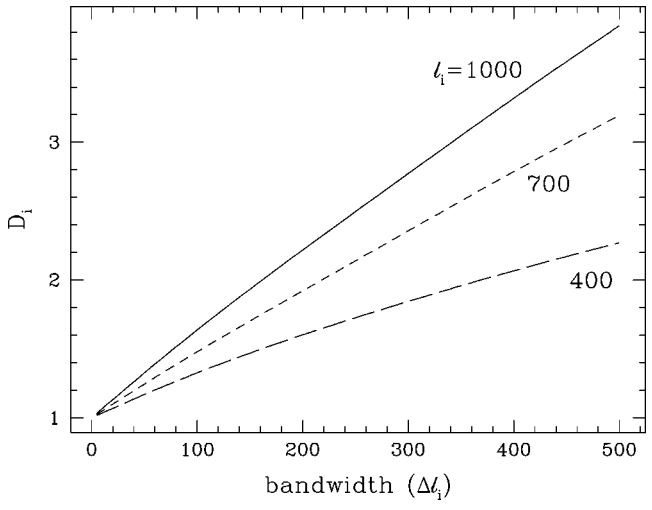

FIG. 3. Band-power variance degradation factor D_i as a function of bandwidth for various choices of the central l_i of the band. The degradation increases with the bandwidth since the Gaussian contribution decreases with the number of wavemodes. The trend illustrates the effect of the non-Gaussian correlation.

IV. SCIENTIFIC IMPACT

The sample covariance of band powers provides the ultimate limitations for measuring the B mode power. However until experiments map the B field at a high signal-to-noise ratio, detector noise will bring the errors closer to Gaussian.

For a given experiment, the covariance matrix of the band powers will include contributions from detector noise and the instrumental beam. The sample covariance matrix of Eq. (14) is replaced with the full covariance matrix

$$C_{ij} = \langle (\hat{\Delta}_i^2 - \Delta_i^2)(\hat{\Delta}_j^2 - \Delta_j^2) \rangle = C_{ij}^G + S_{ij}^N . \tag{20}$$

Let us make the usual approximation that the noise is white and Gaussian. Then [13]

$$C_{ij}^G = \delta_{ij} \frac{2(2\pi)^2}{A\alpha_i^2} \int_{l \in i} d^2 l \left[\frac{l^2}{2\pi} (C_l^{BB} + N_l) \right]^2 ,$$

$$N_l = \left(\frac{\Delta_P}{T} \right)^2 e^{l(l+1)\theta^2_{\mathrm{FWHM}}/8 \ln 2}, \tag{21}$$

where Δ_P^2 is the polarization noise variance in a steradian and θ is the full width at half maximum (FWHM) of the beam in radians.

The impact of the non-Gaussian errors in the power spectrum on the errors in a set of cosmological parameters p_μ can be estimated via the Fisher matrix

$$F_{\mu\nu} = \sum_{ij} \frac{\partial \Delta_i^2}{\partial p_\mu} (\mathbf{C}^{-1})_{ij} \frac{\partial \Delta_j^2}{\partial p_\nu} . \tag{22}$$

The errors on a given parameter $\sigma(p_\mu) \approx (\mathbf{F}^{-1})^{1/2}_{\mu\mu}$. Because constraints on the cosmological parameters of interest, e.g. the neutrino mass and the dark energy properties, will be affected by constraints from the temperature and E power spectra, we defer a full forecast to a future work. Instead we calculate the net degradation in the signal to noise or equiva-

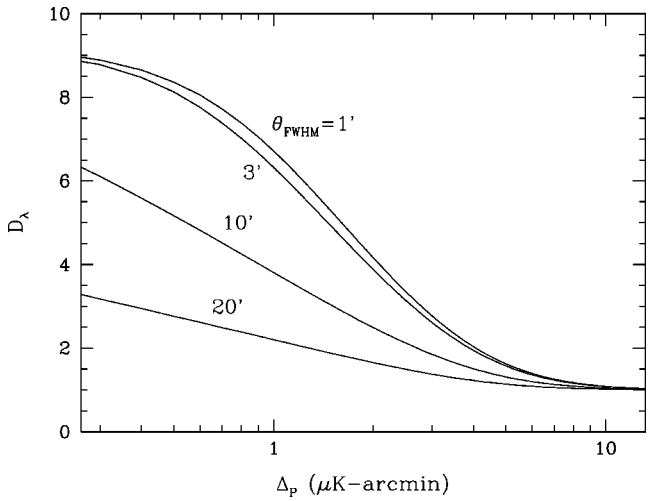

FIG. 4. Non-Gaussian degradation on the amplitude λ of the *B* power spectrum as a function of detector noise for several choices of the FWHM beam $\theta_{\rm FWHM}$. For beams that resolve the *B* power out to the maximum $l_{\rm max}=2000$ ($\theta_{\rm FWHM}\lesssim 3'$), the degradation rapidly rises below $\Delta_P=4\ \mu$K arcmin to the high signal-to-noise, sample variance limit of ~ 10. Increasing $l_{\rm max}$ would further increase the importance of sample variance.

lently the degradation in the errors of the amplitude of the *B* power spectrum given a fixed shape,

$$C_l^{BB}=\lambda C_l^{BB}|_{\rm fid}. \tag{23}$$

For reference, the *B* power spectrum at the peak $l\sim 10^3$ has a sensitivity to the dark energy equation of state *w* of $d\lambda/dw|_{\Omega_m}\approx -0.6$ and to the neutrino mass in eV of $d\lambda/dm_\nu\approx -0.5$. However these sensitivities should not be used for error propagation since they carry an *l* dependence and depend on what is being fixed by the other spectra. For example, the sensitivity to *w* at fixed angular diameter distance to last scattering is $d\lambda/dw|_{D_A}\approx -0.2$. Nonetheless the errors on λ serve as a useful quantification of the overall effect of the non-Gaussianity.

From Eq. (22) the variance in the amplitude becomes

$$\sigma^2(\lambda)=\left(\sum_{ij}\ \Delta_i^2(\mathbf{C}^{-1})_{ij}\Delta_j^2\right)^{-1}. \tag{24}$$

Note that in the Gaussian limit and $\Delta l=1$, the variance becomes

$$\sigma^2(\lambda)|_G=\left[\sum_l\ \frac{2l+1}{2}f_{\rm sky}\left(\frac{C_l^{BB}}{C_l^{BB}+N_l}\right)^2\right]^{-1}. \tag{25}$$

In Fig. 4, we show the degradation factor

$$D_\lambda=\frac{\sigma^2(\lambda)}{\sigma^2(\lambda)|_G} \tag{26}$$

as a function of the noise Δ_P in this $\Delta l=1$ limit. The degradation exceeds a factor of 2 for experiments with $\Delta_P\leqslant 4\ \mu K$ arcmin and a beam $\theta\leqslant 3'$. At this level, the noise within the beam is comparable to the rms of the *B* field [see Eq. (9)]. In the limit of zero noise, the degradation is a factor of ~ 10 for bands out to $l=2000$.

Note that the degradation factor is insensitive to the choice of banding. Computing the variance with the wide bands $\Delta l/l\approx 0.25$ bands of Table I causes a negligible change in comparison to the accuracy of the underlying gradient approximation [see Eq. (5)]. The degradation factor is also insensitive to the area or $f_{\rm sky}$ for a large contiguous region. The degradation can be reduced somewhat by composing the total area out of many smaller patches separated by many degrees on the sky. However this strategy will compromise the separability of *E* and *B* modes. Optimizing a scan strategy against realistic correlated noise, non-Gaussian sample variance and mode leakage will require experiment-specific Monte Carlo simulations.

V. DISCUSSION

The non-Gaussianity of the *B* modes in the lensed CMB polarization substantially degrades the amount of information contained in the *B* mode power spectrum. It both increases the variance of band powers and makes them strongly covary across a wide range in *l* surrounding the peak power. Ultimately it will increase the variance of the amplitude of the power spectrum by an order of magnitude.

As experiments move from the upper limit and first detection stage to using the *B* mode power spectrum to constrain the properties of dark components such as the neutrinos and dark energy, this non-Gaussianity will have to be included in the analysis. By quantifying the sample covariance, we have provided the analytic and numerical tools that will be the basis for such an analysis. The advantage of the Monte Carlo approach is that it can be straightforwardly applied to any estimator of *B* power.

In principle, one can include the sample covariance in an effective χ^2 as is done for power spectrum errors of the temperature field (e.g. Ref. [15]). However since the computation of the covariance is much more costly than the computation of the *B* power spectrum, minimization in a large-dimensional cosmological parameter space is impractical even with Monte Carlo Markov Chain techniques. Since most of the parameters affecting the high redshift universe will be fixed prior to these measurements from the *T* and *E* mode spectra, as a first order correction one can follow the Fisher matrix approach and calculate the effect in a fiducial model. More specifically, given the correlation matrix (R_{ij}) and the relative variance degradation (D_i) in a fiducial model, one can scale the covariance matrix to the model *B* mode spectrum (C_l^{BB}) as calculated from Boltzmann codes. This approach would capture the main scaling of the covariance through the amplitude of the lensing power spectrum. Implementing such a pipeline though is beyond the scope of this paper.

As experiments move from parameter constraints based on power spectra to mapping the lensing potential [6,7], the non-Gaussianity of the polarization becomes the signal and not the noise. The extent to which this ultimate goal will be achievable instrumentally and in the presence of foregrounds [16] awaits the results of the upcoming generation of experiments.

ACKNOWLEDGMENT

We thank T. Okamoto and B. Winstein for useful discussions. K.M.S. and W.H. were supported by NASA Contract No. NAG5-10840, the DOE and the Packard Foundation. M.K. was supported by the NSF and NASA Contract No. NAG-11098. This work was carried out in part at the CfCP under NSF Contract No. PHY-0114422.

[1] M. Zaldarriaga and U. Seljak, Phys. Rev. D **58**, 023003 (1998).
[2] D.N. Spergel *et al.*, Astrophys. J., Suppl. **148**, 175 (2003).
[3] W. Hu, Phys. Rev. D **65**, 023003 (2002).
[4] M. Kaplinghat, L. Knox, and Y.S. Song, Phys. Rev. Lett. **91**, 241301 (2003).
[5] W. Hu, Phys. Rev. D **64**, 083005 (2001).
[6] W. Hu and T. Okamoto, Astrophys. J. **574**, 566 (2002).
[7] C.M. Hirata and U. Seljak, Phys. Rev. D **68**, 043001 (2003).
[8] L. Knox and Y.S. Song, Phys. Rev. Lett. **89**, 011303 (2002).
[9] M. Kesden, A. Cooray, and M. Kamionkowski, Phys. Rev. Lett. **89**, 1304 (2002).
[10] A. Blanchard and J. Schneider, Astron. Astrophys. **184**, 1 (1987).
[11] F. Bernardeau, Astron. Astrophys. **324**, 15 (1997).
[12] W. Hu, Phys. Rev. D **62**, 043007 (2000).
[13] L. Knox, Phys. Rev. D **52**, 4307 (1995).
[14] M. Tegmark, Phys. Rev. D **56**, 4514 (1997).
[15] L. Verde *et al.*, Astrophys. J., Suppl. **148**, 195 (2003).
[16] M. Bowden *et al.*, Mon. Not. R. Astron. Soc. **349**, 321 (2003).