

Lens covariance effects on likelihood analyses of CMB power spectra

Pavel Motloch^{1,2} and Wayne Hu^{1,3}

¹*Kavli Institute for Cosmological Physics, Enrico Fermi Institute,
University of Chicago, Chicago, Illinois 60637, U.S.A*

²*Department of Physics, University of Chicago, Chicago, Illinois 60637, U.S.A*

³*Department of Astronomy & Astrophysics, University of Chicago, Chicago, Illinois 60637, U.S.A*

Non-Gaussian correlations induced in CMB power spectra by gravitational lensing must be included in likelihood analyses for future CMB experiments. We present a simple but accurate likelihood model which includes these correlations and use it for Markov Chain Monte Carlo parameter estimation from simulated lensed CMB maps in the context of Λ CDM and extensions which include the sum of neutrino masses or the dark energy equation of state w . If lensing-induced covariance is not taken into account for a CMB-S4 type experiment, the errors for one combination of parameters in each case would be underestimated by more than a factor of two and lower limits on w could be misestimated substantially. The frequency of falsely ruling out the true model or finding tension with other data sets would also substantially increase. Our analysis also enables a separation of lens and unlensed information from CMB power spectra, which provides for consistency tests of the model and, if combined with other such measurements, a nearly lens-sample-variance free test for systematics and new physics in the unlensed spectrum. This parameterization also leads to a simple effective likelihood that can be used to assist model building in case consistency tests of Λ CDM fail.

I. INTRODUCTION

Measurements of the cosmic microwave background (CMB) have been instrumental in our understanding of the composition and evolution of the Universe (see e.g. [1]). Starting with the initial detection through cross correlation of the quadratic reconstruction of the lensing potential [2] with galaxy surveys [3], gravitational lensing of the CMB by the large scale structure of the Universe (see [4] for a review) has emerged as a new source of cosmological information [5–13]. For currently operating and future CMB experiments [14–17], gravitational lensing will be instrumental in breaking parameter degeneracies affecting low redshift physics, such as sum of the neutrino masses and properties of the dark energy.

On the other hand, lensing also acts as a source of lens-sample-variance noise for the detection of inflationary gravitational waves and light relic particles from CMB power spectra (see e.g. [17]). Lens sample variance can be removed by reconstructing the lensing potential assuming that the power spectra and higher point moments are lensed by the same potential. Recently, Ref. [18] proposed a consistency test between the two that can be used to protect against systematic errors and incorrect assumptions on either side.

The covariance between CMB observables induced by lensing also complicates their analysis. This covariance was first studied between the lensed temperature and polarization power spectra [19, 20], yet all analyses of CMB experiments to date have omitted it in the likelihood function (e.g. [1]). Once polarization measurements approach the sample variance limit, this covariance must be included in the analysis [20] as well as the correlation between power spectra and the reconstructed lensing potential [18, 21, 22].

Most previous studies have employed Fisher forecasts

to estimate lens-covariance effects on parameter constraints rather than a direct likelihood exploration, with the exception of [22] who briefly investigate their impact on neutrino mass constraints in Λ CDM. The main goal of this work is filling this gap with a more comprehensive study of an analysis pipeline from simulated lensed CMB maps to parameter posterior probabilities, including the principal component parameterization of the lensing potential introduced in Ref. [18].

This paper is organized as follows. In § II we present our model for the likelihood function of full-sky CMB power spectra data which accounts for lensing-induced covariance between multipole moments. We briefly review the analytical model for the lensing-induced covariance upon which it is based [20] and define the fiducial cosmology and experimental configuration investigated in this work. Then in § III we look at effects the non-Gaussian covariance has on cosmological parameter constraints and point out that neglecting this covariance could have a possibly serious impact on concordance studies. In § IV we introduce a principal component parameterization of lensed CMB power spectra, describe results of our MCMC analyses of properties of these parameters and use them to explain the effects of lensing-induced covariance on cosmological parameters. In the same section we then detail how measurements of these effective parameters can be used to form consistency tests for the data and assist model building.

In the Appendices we give more details on how we modify code `Lenspix` which simulates lensed CMB data to achieve high accuracy (Appendix A) and how we determined necessary number of lensing principal components (Appendix B).

TABLE I. Fiducial parameters used in the analysis

| Parameter | Fiducial value |
|----------------|------------------------|
| h | 0.675 |
| $\Omega_c h^2$ | 0.1197 |
| $\Omega_b h^2$ | 0.0222 |
| n_s | 0.9655 |
| A_s | 2.196×10^{-9} |
| τ | 0.06 |
| $\sum m_\nu$ | 60 meV |

II. ANALYSIS DETAILS

In this section we first describe our fiducial cosmology and experimental configuration. Then we comment on analytic model used for the covariance of lensed CMB power spectra, details of our CMB simulations and Markov Chain Monte Carlo analysis and in a separate section we describe our model for the data likelihood.

A. Fiducial cosmology and experimental setup

In this work we investigate a simplified experimental setup of a full sky experiment with specifications inspired by CMB Stage 4 [17]. Throughout the paper we use capital letters X, Y, W, Z to represent either CMB temperature or polarization field, i.e. an element from $\{T, E, B\}$. In a given cosmological model, we also abbreviate the set of all cosmological parameters as θ_A . We consider the information on θ_A provided by the CMB power spectra C_ℓ^{XY} .

For the fiducial cosmology we take a flat Λ CDM model with minimal neutrino mass. For the Λ CDM parameters we take $\omega_b = \Omega_b h^2$, the physical baryon density; $\omega_c = \Omega_c h^2$, the physical cold dark matter density; n_s , the tilt of the scalar power spectrum; A_s , its amplitude; and τ , the optical depth to recombination. We choose θ_* , the angular scale of the sound horizon at recombination, as opposed to the Hubble constant h , as the sixth independent parameter. We also assume that tensor modes are negligible so that there is no unlensed B mode. Values of the cosmological parameters for the fiducial model used in this work are summarized in Table I.

For noise in temperature and polarizations, we assume a noise spectra [23]

$$N_\ell^{XY} = \Delta_{XY}^2 e^{\ell(\ell+1)\theta_{\text{FWHM}}^2/8\ln 2}, \quad (1)$$

where Δ_{XY} is the instrumental noise (in μK -radian) and θ_{FWHM} is the beam size (in radians). We consider a $1'$ beam, $\Delta_{TT} = 1 \mu\text{K}'$, $\Delta_{EE} = \Delta_{BB} = 1.4 \mu\text{K}'$, and $\Delta_{TE} = \Delta_{TB} = \Delta_{EB} = 0$ and use measurements in the multipole range $\ell = 2 - 3000$.

B. Data covariance

Because all CMB temperature and polarization anisotropies are lensed by the same realization of the lensing potential, the lensed CMB power spectra data are correlated across multipoles. The ensuing covariance can be well described by a simple analytical model which has been tested on simulations [20]. Recently, this model has been extended to capture correlations of the lensed CMB power spectra with the lensing potential $C_\ell^{\hat{\phi}\hat{\phi}}$ reconstructed by a quadratic estimator [18, 21, 22]. Here we describe salient features of this model, more in-depth discussion and graphical representation of the resulting covariances can be found in the references.

In this model the correlation matrix is split into ‘‘Gaussian part’’ \mathcal{G} that is diagonal in multipole space and \mathcal{N} which describes non-Gaussian correlations between multipoles,

$$\text{Cov}_{\ell\ell'}^{XY,WZ} = \mathcal{G}_{\ell\ell'}^{XY,WZ} + \mathcal{N}_{\ell\ell'}^{XY,WZ}. \quad (2)$$

The Gaussian part is modeled after the covariance of Gaussian random fields as

$$\mathcal{G}_{\ell\ell'}^{XY,WZ} = \frac{\delta_{\ell\ell'}}{2\ell+1} [C_{\text{exp},\ell}^{XW} C_{\text{exp},\ell}^{YZ} + C_{\text{exp},\ell}^{XZ} C_{\text{exp},\ell}^{YW}], \quad (3)$$

where the expectation value of the experimentally measured lensed CMB power spectra C_{exp}^{XY} includes the noise power spectrum N_ℓ^{XY}

$$C_{\text{exp},\ell}^{XY} = C_\ell^{XY} + N_\ell^{XY}. \quad (4)$$

Even if we assume that the unlensed CMB fields \tilde{X} and ϕ are Gaussian, the lensed CMB fields X are not. In our model, we take two non-Gaussian terms to compose the full covariance,

$$\mathcal{N}_{\ell\ell'}^{XY,WZ} = \mathcal{N}_{\ell\ell'}^{(\phi)XY,WZ} + \mathcal{N}_{\ell\ell'}^{(E)XY,WZ}. \quad (5)$$

The first term is

$$\mathcal{N}_{\ell\ell'}^{(\phi)XY,WZ} = \sum_L \frac{\partial C_\ell^{XY}}{\partial C_L^{\phi\phi}} \text{Cov}_{LL}^{\phi\phi} \frac{\partial C_{\ell'}^{WZ}}{\partial C_L^{\phi\phi}} \quad (6)$$

and corresponds to correlations induced by the common set of gravitational lenses. The power spectra derivatives are in practice calculated using a two point central difference scheme from results obtained using CAMB^{*1} [24].

Sample variance of the unlensed $\tilde{E}\tilde{E}$ power spectrum and its coherent propagation into the lensed power spectra through gravitational lensing produces similar but typically weaker effects. Following [20] we include this contribution only for $\text{Cov}_{\ell\ell'}^{XY,BB}$ with

$$\mathcal{N}_{\ell,\ell'}^{(E)XY,BB} = \sum_L \frac{\partial C_\ell^{XY}}{\partial C_L^{\tilde{X}\tilde{Y}}} \text{Cov}_{L,L}^{\tilde{X}\tilde{Y},\tilde{E}\tilde{E}} \frac{\partial C_{\ell'}^{BB}}{\partial C_L^{\tilde{E}\tilde{E}}}. \quad (7)$$

Other sample covariance effects from unlensed fields on XY are negligible in comparison [20].

^{*1} <http://camb.info>

C. Simulated data and their analysis

To simulate lensed CMB data we use the publicly available code `Lenspix`^{*2} [25] with unlensed CMB power spectra calculated by `CAMB`. We modified the code to lower its memory demands and speed up the calculation, see Appendix A for details on these modifications. Once the lensed CMB maps are generated, we add normally distributed^{*3} instrumental noise and calculate power spectra to form a simulated data set \hat{C}_ℓ^{XY} .

In most of the analyses in this paper we investigate the simulated CMB power spectra using the Markov Chain Monte Carlo (MCMC) code `CosmoMC`^{*4} [26], which for a given realization of the data samples the posterior probability in the space of cosmological parameters. We assume uniform priors in the cosmological parameters and use the likelihood described in the next section. In the MCMC runs we sample the posterior until the Gelman-Rubin statistic $R - 1$ [27] drops below 0.01.

To avoid biasing results, we calculate the unlensed fiducial spectra which enter `Lenspix` simulations with the same precision settings which is later used in `CosmoMC`. We checked that increasing precision with which the lensing operation in `CosmoMC` is calculated (increasing `accuracy_boost` in the lensing routine) does not have any effect on the resulting parameter constraints.

D. Likelihood

An accurate likelihood for CMB power spectra data \hat{C}_ℓ^{XY} has to capture both lensing-induced covariance $\text{Cov}_{\ell,\ell'}^{XY,WZ}$ and the non-normal distribution of the power spectra at low multipoles. Here we illustrate these effects with simulated data and then describe our model for the likelihood.

Using 2000 `Lenspix` simulations, it is possible to illustrate that the lensed CMB data are indeed correlated. Because this number of simulations is insufficient to show correlation of individual power spectra multipoles, we look at correlation between band powers

$$P_{\ell_1,\ell_2}^{XY} = \frac{1}{\ell_2 - \ell_1} \sum_{\ell=\ell_1}^{\ell_2} \frac{\ell(\ell+1)\Delta C_\ell^{XY}}{2\pi}, \quad (8)$$

where

$$\Delta C_\ell^{XY} = \hat{C}_\ell^{XY} - C_{\text{exp},\ell}^{XY} \quad (9)$$

is the deviation of the experimentally measured power spectrum from its expectation value. As an example, in




FIG. 1. Light gray circles show the correlated values of the binned power spectra $P_{1000,1300}^{BB}$ and $P_{1500,1800}^{BB}$ as determined from 2000 `Lenspix` simulations. Blue lines encompass regions of 68% and 95% confidence determined from these simulations. For comparison, dashed red lines show the same confidence intervals based on our theoretical model for covariance.

Figure 1 we plot the distribution of two BB band powers as determined from our simulations, together with theoretical curves showing 68% and 95% confidence intervals derived from our model for the covariance $\text{Cov}_{\ell,\ell'}^{XY,WZ}$. We see that the data are indeed strongly correlated and that the model describes this correlation well.

Fortunately for the likelihood construction, non-Gaussian covariances $\mathcal{N}_{\ell\ell'}^{XY,WZ}$ of the low ℓ data are considerably smaller than the corresponding Gaussian part of the covariance. Indeed, the largest correlation coefficient

$$R_{\ell\ell'}^{XY,WZ} = \frac{\mathcal{N}_{\ell\ell'}^{XY,WZ}}{\sqrt{\text{Cov}_{\ell\ell}^{XY,XY} \text{Cov}_{\ell'\ell'}^{WZ,WZ}}} \quad (10)$$

with $\ell < 30$ is 8×10^{-4} . We shall therefore neglect lensing-induced covariances in the likelihood for the large scale data.

As pointed out above, the low ℓ data are not normally distributed. For example, in Fig. 2 we plot the distribution of B -mode power spectra including noise for $\ell = 2, 30$ obtained in the same `Lenspix` simulations. In the same plot we show a normal distribution centered on the expected mean of the data, with standard deviation given by the Gaussian expectation (3)

$$\sigma_\ell^{BB} = \sqrt{\frac{2}{2\ell+1}} C_{\text{exp},\ell}^{BB}. \quad (11)$$

It is clear that for $\ell = 2$ the normal distribution is a poor description of the data. Instead, as expected, $\chi_{2\ell+1}^2$ distribution scaled by σ_ℓ^{BB} fits the data well. In the case of large scale B -modes this reflects the fact that they

^{*2} <https://github.com/cmbant/lenspix>

^{*3} To avoid confusion between Gaussian covariance and Gaussian distributions, we use “normal” for the latter.

^{*4} <https://github.com/cmbant/CosmoMC>

get most of their power from \tilde{E}, ϕ modes at ℓ of several hundred. Each coefficient in the spherical harmonic expansion of the B map is then a combination of many random fields and thus approximately normally distributed, which leads to χ^2 distributed power spectra. For low ℓ TT, TE and EE , the distributions just mirror the unlensed CMB fields, due to negligible effects of lensing on these fields. Above $\ell \sim 30$ the normal distribution becomes a good description for both the χ^2 distribution and the data.

Based on these considerations, our model for the likelihood treats \hat{C}_ℓ^{XY} from the largest scales ($\ell < \ell_{\text{break}}$) and from smaller scales separately and independently. The choice of the division point is somewhat arbitrary; in this work we use $\ell_{\text{break}} = 30$.

Below ℓ_{break} we neglect lensing-induced covariance $\mathcal{N}_{\ell, \ell'}^{XY, WZ}$ and assume there is no correlation between B modes and T, E modes. Data with different multipoles ℓ, ℓ' then decouple. The likelihood of measuring the data vector \hat{C}_ℓ^{XY} (including instrumental noise) when the expected power spectra are $C_{\text{exp}, \ell}^{XY}$ is then a product of inverse Wishart distributions

$$\begin{aligned} \mathcal{L}_{\ell < \ell_{\text{break}}} &\propto \prod_{\ell=2}^{\ell_{\text{break}}-1} |C_{\text{exp}, \ell}^{XY}|^{-(2\ell+1)/2} \\ &\times \exp\left(-\frac{2\ell+1}{2} \sum_{X, Y} (C_{\text{exp}, \ell}^{XY})^{-1} \hat{C}_\ell^{XY}\right). \end{aligned} \quad (12)$$

Here $|\cdot|$ is a determinant of \cdot viewed as a matrix, here in the X, Y space.

For $\ell \geq \ell_{\text{break}}$ we neglect the non-normality of the distribution of each multipole and instead model the lensing-induced covariance between multipoles:

$$\begin{aligned} \mathcal{L}_{\ell \geq \ell_{\text{break}}} &\propto |\text{Cov}|^{-1/2} \\ &\exp\left[-\frac{1}{2} \sum_{\substack{\ell, \ell' \geq \ell_{\text{break}} \\ XY, WZ}} \Delta C_\ell^{XY} (\text{Cov}_{\ell, \ell'}^{XY, WZ})^{-1} \Delta C_{\ell'}^{WZ}\right]. \end{aligned} \quad (13)$$

In all analyses in this paper we neglect the dependence of the covariance matrix on the cosmological parameters by evaluating it at the fixed fiducial model of Tab. I.

To assess the impact of non-Gaussian covariance, we also investigate the likelihood $\mathcal{L}_{\text{g}, \ell \geq \ell_{\text{break}}}$ in which the non-Gaussian covariance $\text{Cov}_{\ell, \ell'}$ is replaced by the Gaussian covariance $\mathcal{G}_{\ell, \ell'}$.

By joining the large and small scale portions independently, we then form the total likelihood for the data

$$\ln \mathcal{L} = \ln \mathcal{L}_{\ell < \ell_{\text{break}}} + \ln \mathcal{L}_{\ell \geq \ell_{\text{break}}} \quad (14)$$

and

$$\ln \mathcal{L}_{\text{g}} = \ln \mathcal{L}_{\ell < \ell_{\text{break}}} + \ln \mathcal{L}_{\text{g}, \ell \geq \ell_{\text{break}}}, \quad (15)$$

up to irrelevant additive constants.

By comparing analyses based on these two likelihoods we are able to gauge the impact of non-Gaussian covariance on parameter constraints. The lensing-induced co-

variance is an additional source of correlated noise so constraints based on the likelihood with Gaussian covariance are typically too optimistic.

III. COSMOLOGICAL PARAMETER ESTIMATION

In this section we investigate how neglecting lensing-induced covariance affects constraints on cosmological parameters. We focus here on the six parameter flat Λ CDM cosmological model and two extensions where either the sum of the masses of neutrinos $\sum m_\nu$ (Λ CDM+ $\sum m_\nu$) or the dark energy equation of state parameter w (Λ CDM+ w) is allowed to vary. Similar studies on the impact of non-Gaussian covariances have been previously performed mainly using the Fisher approximation [18–20, 22]. We find that in at least one case (Λ CDM+ w) the Fisher approximation significantly underestimates the impact of lensing-induced covariance.

For Λ CDM we also illustrate how neglecting lensing-induced covariance leads to a significant increase in the fraction of realizations in which the fiducial model parameters are excluded at 95% confidence, which is potentially important for concordance studies.

A. Λ CDM

Neglecting lensing-induced covariances for a typical simulated CMB dataset affects constraints on Λ CDM parameters as shown in Fig. 3. Shifts in the best fit parameter values of the base parameters of Tab. I are typically not very significant; the major effect of including lensing covariances is a weakening of the best constrained directions between degenerate parameters, most notably that between $\Omega_c h^2$ and A_s . We comment on the origin of this effect in §IV C.

Because of marginalization of other parameters, the two parameter posteriors in Fig. 3 hide some of the effects of the non-Gaussian covariance. To uncover the maximal possible effect on a single quantity, we construct the linear combination of the cosmological parameters

$$M = \sum_A \mathcal{K}_A (\theta_A - \theta_A^{\text{fid}}) \quad (16)$$

that maximizes the ratio of Gaussian to non-Gaussian errors. Here $\mathcal{K}_A = \{5.8, -13.4, 18.4, -1.1, -2.6, 3.1\}$ for the parameter ordering $\{100\theta_*, \Omega_c h^2, \Omega_b h^2, n_s, \ln A_s, \tau\}$. As we discuss in greater detail in §IV C, M can be interpreted as a combination of cosmological parameters which mainly changes the lensing potential, especially at low ℓ .

Dashed red lines in Fig. 4 show posterior probabilities for M , as determined from a single MCMC run based on \mathcal{L} (left) and \mathcal{L}_{g} (right). The same simulated CMB sky as in Fig. 3 is used, the maximum has been shifted to zero for future convenience, and the y -axis units are arbitrary.




FIG. 2. Distribution of $\hat{C}_2^{BB}, \hat{C}_{30}^{BB}$ obtained from 2000 Lenspix simulations compared against a χ^2 distribution with the Gaussian variance $(\sigma_\ell^{BB})^2$ (solid) and a normal distribution with the same variance centered on the expected mean of the data (dashed). The normal distribution becomes a good approximation for $\ell \gtrsim 30$.

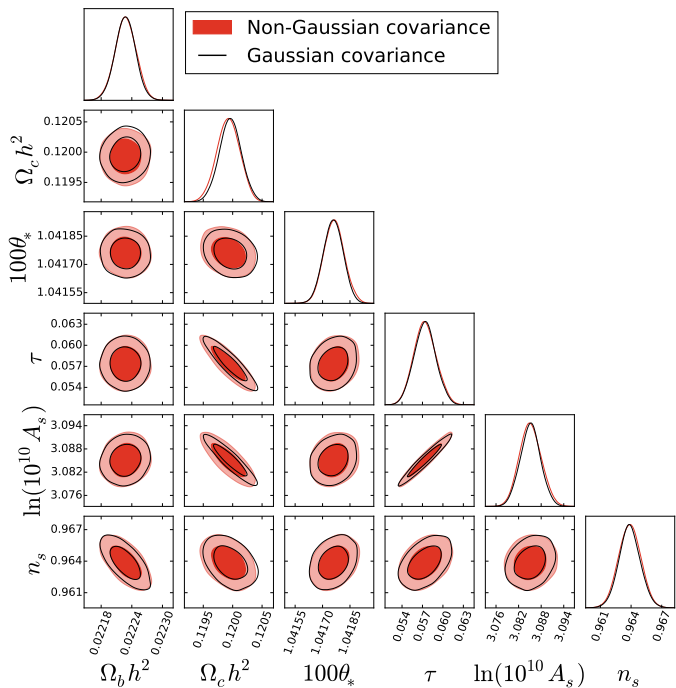


FIG. 3. Comparison of MCMC constraints on Λ CDM parameters with analysis based on Gaussian (black curves) and non-Gaussian covariance (red shaded). Here and throughout contours enclose regions of 68% and 95% confidence intervals unless otherwise specified.

The standard deviations of these two posteriors are $\sigma_M^{\text{ng}} = 2.1 \times 10^{-3}$ and $\sigma_M^{\text{g}} = 1.0 \times 10^{-3}$, both within 3% of the Fisher forecast prediction displayed in Fig. 4 by a solid blue line. The analysis based on a Gaussian likelihood \mathcal{L}_g therefore underestimates the errors of M by over a factor of 2.

While the impact of this direction is hidden in the marginalized errors of the base parameters, it reveals itself in an increased frequency of Type 1 errors: falsely rejecting the true model. This effect would be particularly problematic for concordance studies searching for tensions between various cosmological datasets.

Since a full study of Type 1 errors in thousands of data realizations is computationally expensive, we illustrate this problem by analytically approximating the best fit values of the parameters, including M , for each realization. We assume the data are sufficiently close to the fiducial model that it is possible to approximate the C_ℓ^{XY} as linear in the parameter deviations $\theta_A - \theta_A^{\text{fid}}$. Neglecting for the moment complications arising from presence of the low ℓ data by assuming all the data are distributed according to a multivariate normal distribution with covariance (2), we obtain the maximum likelihood or best fit estimate for a cosmological parameter θ_A^{bf} as^{*5}

$$\begin{aligned} \Delta\theta_A &\equiv \theta_A^{\text{bf}} - \theta_A^{\text{fid}} \\ &= \sum_{B, i, j} F_{AB}^{-1} \frac{\partial D_i}{\partial \theta_B} (\text{Cov}^{-1})_{ij} (\hat{D}_j - D_j^{\text{fid}}). \end{aligned} \quad (17)$$

Here the power spectrum data are indexed as $D_i = C_\ell^{XY}$ with i running over all XY, ℓ elements. The Fisher information matrix,

$$F_{AB} = \sum_{ij} \frac{\partial D_i}{\partial \theta_A} (\text{Cov})_{ij}^{-1} \frac{\partial D_j}{\partial \theta_B}, \quad (18)$$

^{*5} Numerical maximization of the likelihood in ten simulations provides an average shift in M below $0.08 \sigma_{M, \text{bf}}$ with respect to the analytic formula (17) for either covariance. The analytic treatment is thus sufficiently accurate for our purposes.

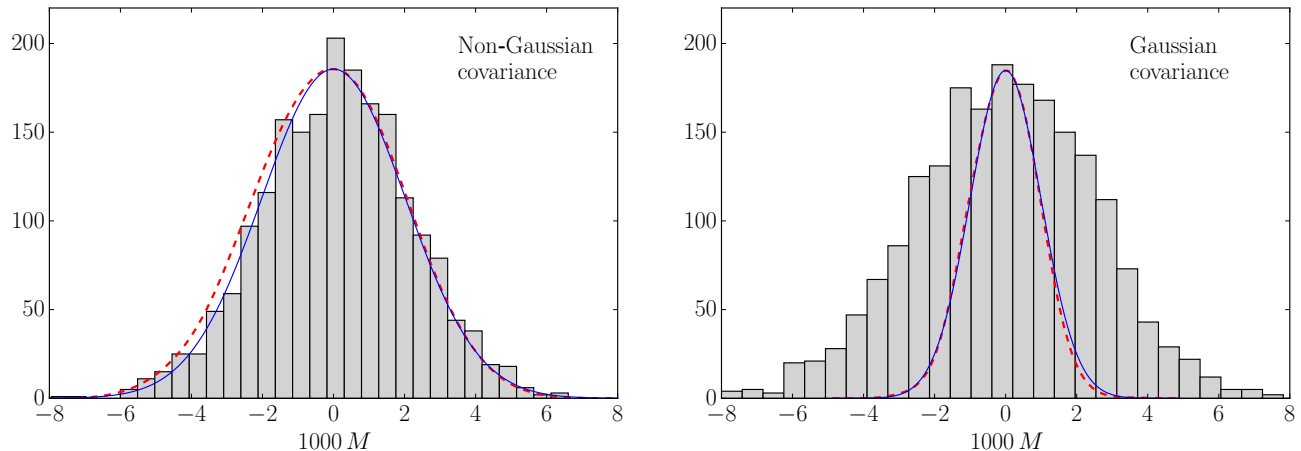


FIG. 4. Histograms showing best fit values of M determined from 2000 lensed CMB skies using an analytic approach (see text) for analysis based on non-Gaussian (left) and Gaussian covariance (right). The dashed red curves show the 1D posterior probability for M , as determined from a single CMB realization, shifted to zero for better comparison. Despite the Fisher forecast (solid blue curves) agreeing in both cases, the posterior probability does not reflect the much wider best-fit distribution in the Gaussian case.

and parameter derivatives are evaluated around the fiducial model.

With the best fit values of cosmological parameters we can directly calculate the best fit values of M for both the non-Gaussian and Gaussian analysis. The two differ only in the choice of covariance matrix in Eqs. (17) and (18).

Best fit values of M determined from 2000 simulated CMB skies are shown in histograms in Fig. 4 for the non-Gaussian (left) and Gaussian (right) covariance analyses. In the non-Gaussian case, this distribution has a standard deviation $\sigma_{M,\text{bf}}^{\text{ng}} = 2.1 \times 10^{-3}$ which is in excellent agreement with the prediction from the posterior σ_M^{ng} determined from the MCMC analysis of single realization as well as the Fisher approximation. On the other hand, in the Gaussian analysis the best fit values of M scatter with standard deviation $\sigma_{M,\text{bf}}^{\text{g}} = 2.5 \times 10^{-3}$, which is 2.5 times the width of σ_M^{g} despite the latter agreeing with its Fisher approximation. This mismatch can lead to Type 1 errors in cases where the best fit M fluctuates away from the fiducial value zero.

Notice that the best fit distribution is wider in the Gaussian than non-Gaussian case by a factor of ~ 1.2 which further exacerbates the probability of Type 1 errors. This is not surprising, Eq. (17) is a minimum variance estimator only if the assumed model of the covariances is correct which it is not in the Gaussian case.

To quantify the probability of Type 1 errors considering all parameters that specify the Λ CDM model, we can also compute χ^2 between the best fit and the true fiducial model assuming the errors from the posterior

$$\chi^2(\Delta\theta_A) = \sum_{AB} \Delta\theta_A \text{Cov}_{AB}^{-1} \Delta\theta_B \approx \sum_{AB} \Delta\theta_A F_{AB} \Delta\theta_B. \quad (19)$$

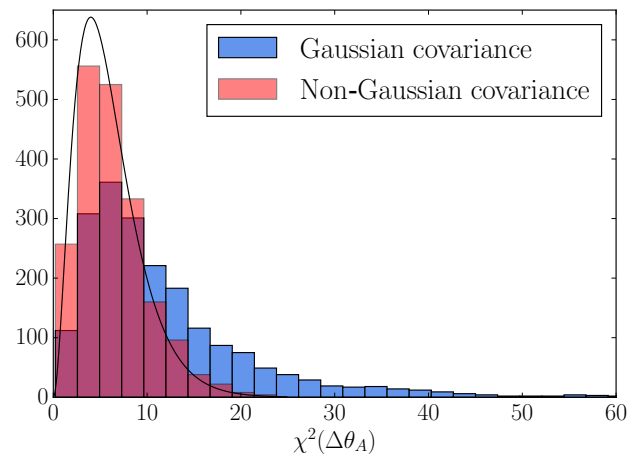


FIG. 5. Histogram of $\chi^2(\Delta\theta_A)$ for the parameter deviations of the best fit from the true model (19), as determined from our simulations with Gaussian (blue) and non-Gaussian (red) covariance. For comparison, the solid line is proportional to probability density function of χ_6^2 . The long tail in the Gaussian case leads to anomalously frequent Type 1 errors where the true model is rejected at high confidence (see text).

In order to estimate χ^2 for each of the 2000 lensed CMB simulations, we again use the Fisher matrix as an approximation to the inverse covariance.^{*6} The variable

^{*6} Using the ten simulations in which we find actual best fit Λ CDM parameters and use actual posterior covariance for the parameters, we can estimate what is the error from using in (19) the an-

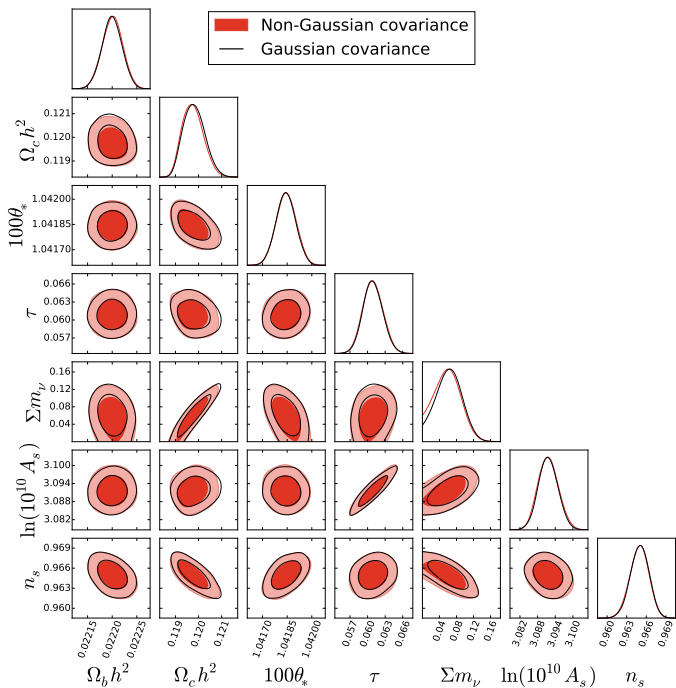


FIG. 6. Comparison of MCMC constraints on $\Lambda\text{CDM} + \sum m_\nu$ parameters with analysis based on Gaussian (black curves) and non-Gaussian covariance (red shaded).

$\chi^2(\Delta\theta_A)$ should be χ_6^2 distributed, where 6 is the number of cosmological parameters in ΛCDM .

Histograms of $\chi^2(\Delta\theta_A)$ for the Gaussian and non-Gaussian analysis are compared in Fig. 5 to the theoretical expectation. It is clear that in the Gaussian analysis, the misestimate of the parameter covariance as well as the suboptimal estimate of the best fit causes a strong disagreement with the expected χ_6^2 distribution.

For example, when the analysis is based on the Gaussian covariance, more than 30% of the simulations show $\chi^2(\Delta\theta_A) > 12.59$; for χ_6^2 this value is exceeded only in 5% of the cases. As pointed out above, this can be potentially dangerous for concordance studies. The non-Gaussian covariance leads to much better agreement ($\sim 6.9\%$ of simulations have $\chi^2(\Delta\theta_A) > 12.59$) and moreover there is no long tail to very high $\chi^2(\theta_A)$.

B. $\Lambda\text{CDM} + \sum m_\nu$

In this section we release the neutrino mass from its fiducial value and investigate a seven-parameter extension of ΛCDM . In this case, the two parameter contour

alytic estimator for $\Delta\theta_A$ and Fisher matrix. The average change in $\chi^2(\theta_A)$ we observe is 0.7, sufficiently smaller than the width of the final distributions. Most of this difference comes from the analytic estimator (17).

plots, shown in Fig. 6, show much smaller effects of non-Gaussian covariance than in ΛCDM , though the impact is visible in the lower limit for $\sum m_\nu$. Ref. [22] also found only small effects for this case.

However, this does not mean non-Gaussian covariance can be neglected in this case, only that its effects are hidden by marginalizations. As for ΛCDM , we can form a combination of the cosmological parameters which is predicted by the Fisher forecast to show the largest effect of non-Gaussian covariance,

$$M^\nu = \sum_A \mathcal{K}_A^\nu (\theta_A - \theta_A^{\text{fid}}) \quad (20)$$

with $\mathcal{K}_A^\nu = \{5.7, -13.3, 18.4, -1.1, -2.6, 3.1, 0.29 \text{ eV}^{-1}\}$, for the ordering $\{100\theta_s, \Omega_c h^2, \Omega_b h^2, n_s, \ln A_s, \tau, \sum m_\nu\}$. In Fig. 7 we can see 1D posterior probabilities for M^ν from MCMC analyses based on Gaussian and non-Gaussian covariance. It is clear M^ν constraints are nearly as strongly affected by the lensing-induced covariance as M for ΛCDM , with standard deviation degrading from 1.1×10^{-3} to 2.2×10^{-3} , even though this effect does not show up in any pair of base parameters. Likewise, the Gaussian analysis is prone to Type 1 errors as in the ΛCDM case.

The standard deviations of M^ν quoted in the previous paragraph differ somewhat from their Fisher forecasts: 1.0×10^{-3} and 2.1×10^{-3} for the Gaussian and non-Gaussian cases respectively. This occurs in part because the posterior is non-normal due to the presence of the physicality prior $\sum m_\nu > 0$ which hides some of the non-Gaussian covariance effects (cf. Fig. 6). In cases where neutrino mass is detected with high significance and effects of the prior boundary is smaller, M^ν constraints are in good agreement with the Fisher prediction.

C. $\Lambda\text{CDM} + w$

In the model where we allow the dark energy parameter of state to vary, the effect of non-Gaussian covariance is more pronounced and clearly visible already on posterior probability distribution for w , see Fig. 8. The two analyses, based on \mathcal{L} and \mathcal{L}_g , strongly disagree in the low w tail; by neglecting the covariance induced by the gravitational lensing one would wrongly rule out low values of w . For example, for this particular realization lower 95% confidence limits (two sided) for the non-Gaussian and Gaussian likelihoods are -1.55 and -1.37 respectively. In §IV C we look deeper into this behavior.

The impact of non-Gaussian covariance depends to some extent on best fit value of w ; simulations with lower best fit value of w typically show larger effects of non-Gaussian covariance. To illustrate this, in Fig. 9 we show marginalized constraints on w for six different simulations, both with Gaussian and non-Gaussian likelihood.

In the Fisher approximation, $\Lambda\text{CDM} + w$ was investigated for essentially the same experimental configuration in [19]. It was found that non-Gaussian covariances




FIG. 7. Effect of the non-Gaussian covariance on constraints on a parameter combination M^ν (20) within Λ CDM+ $\sum m_\nu$; the combination was chosen to maximize this effect. Solid line shows MCMC constraints with non-Gaussian covariance, dashed line with the Gaussian covariance.

should increase errors on w by only about 24%, which is considerably less than what we uncovered in the full MCMC analysis. In this case, the local approximation thus gives misleading results, due to the non-normal posterior which the Fisher approximation can not faithfully capture. The origin of this behavior can be traced back to how dark energy affects lensing. As w decreases, its effect on the lensing potential quickly diminishes, as dark energy ceases to be important at redshifts where the lensing kernel peaks. Because in this case the parameter combination constrained by lensing is significantly changing throughout the allowed parameter posterior, the Fisher analysis fails to capture the full significance of the non-Gaussian covariance.

The maximally impacted linear combination of base parameters also shows an enhanced non-Gaussian effect compared with Λ CDM and Λ CDM+ $\sum m_\nu$. It reads

$$M^w = \sum_A \mathcal{K}_A^w (\theta_A - \theta_A^{\text{fid}}), \quad (21)$$

where $\mathcal{K}_A^w = \{5.7, -12.8, 18.0, -1.0, -2.5, 2.9, 0.087\}$ for the ordering $\{100\theta_*, \Omega_c h^2, \Omega_b h^2, n_s, \ln A_s, \tau, w\}$. Posterior probabilities for M^w from MCMC analyses based on Gaussian and non-Gaussian covariance are shown in Fig. 10. In this case, the standard deviations $\sigma_{M^w, \text{Fisher}}^{\text{ng}} = 2.3 \times 10^{-3}$ and $\sigma_{M^w, \text{Fisher}}^{\text{g}} = 1.0 \times 10^{-3}$ derived using a Fisher approximation show that the relative impact of non-Gaussian covariance is in reasonable agreement with the MCMC results $\sigma_{M^w}^{\text{ng}} = 7.6 \times 10^{-3}$ and $\sigma_{M^w}^{\text{g}} = 3.4 \times 10^{-3}$, but the overall scale is still strongly underestimated by the Fisher analysis as is the extent of the lower tail. These mismatches are expected for the same reason that they appear in w alone, namely due to parameter nonlinearity within the allowed region. Likewise,




FIG. 8. Comparison of MCMC constraints on Λ CDM+ w parameters with analysis based on Gaussian (black curves) and non-Gaussian covariance (red shaded). The impact of non-Gaussian covariance is clearly apparent in constraints involving w .

M^w defined by (21) only captures the parameter combination which is the most affected by the lensing-induced covariance locally at the fiducial model parameters, not necessarily globally (see §IV C).

IV. LENS PRINCIPAL COMPONENTS

The temperature and polarization power spectra contain more information about the lensing potential than just its amplitude; this information can be faithfully captured in terms of a few principal components (PCs) [18, 19]. This model-independent lensing information can be utilized in a variety of ways. By isolating the lensing information, we can more directly diagnose when lensing-induced covariance is important for parameter constraints. Lensing PCs also enable us to check the internal consistency of the data with a given model, by comparing the constraints on parameters derived from the lensing potential with those from the unlensed power spectra. Comparing $C_\ell^{\phi\phi}$ constraints from the power spectra with those from the reconstructed lensing potential and external measurements provides yet another powerful consistency check. Finally, PCs enable model building beyond the currently considered model classes should these consistency tests fail.

In this section we conduct an MCMC study of the lensing PCs $\Theta^{(i)}$ as defined by the Fisher approximation




FIG. 9. Comparison of MCMC constraints on w for our experimental setup; each panel represents a different Λ CDM simulation of the CMB sky with that of Fig. 8 in the top left corner. The Gaussian analysis (dashed) increasingly deviates from the non-Gaussian analysis (solid) as the maximum likelihood value of w decreases.




FIG. 10. Effect of the non-Gaussian covariance on constraints on a parameter combination M^w (21) within Λ CDM+ w ; the combination was chosen to maximize this effect locally around the fiducial model. Solid line shows MCMC constraints with non-Gaussian covariance, dashed line with the Gaussian covariance.

[18, 19] together with variables $\tilde{\theta}_A$ that parametrize the unlensed power spectrum. Since the PCs are defined under the Fisher approximation, we first verify that in the MCMC analysis the PCs remain unbiased and weakly correlated – both mutually and with $\tilde{\theta}_A$. Then we point out that almost the whole effect of the non-Gaussian covariance is manifested in the first two PCs $\Theta^{(1)}, \Theta^{(2)}$ and use this knowledge to explain effects of non-Gaussian covariances on parameter constraints seen in the previous




FIG. 11. Five principal components $K_\ell^{(i)}$ of the lensing potential best measured by the lensed power spectra.

section.

We end this section with a discussion of possible consistency checks using these principal components and a suggestion of how to use them to compress most of the information contained in the lensed CMB power spectra into a simple normal likelihood, which can be used to quickly determine approximate constraints on a wider class of models than explicitly analyzed here.

A. Parameterizing lens and unlensed power spectra

For a particular experiment and fiducial model, a Fisher forecast can determine which principal components $\Theta^{(i)}$ of $C_\ell^{\phi\phi}$ will be the best measured by the XY power spectra. These PCs can be ordered by increasing variance so that only the handful of best measured components need be included in the actual analysis. In Ref. [18], the hierarchy of $\Theta^{(i)}$ was found for the experimental configuration considered here. These principal components are defined as

$$\Theta^{(i)} = \sum_\ell K_\ell^{(i)} (\ln C_\ell^{\phi\phi} - \ln C_\ell^{\phi\phi, \text{fid}}), \quad (22)$$

where $C_\ell^{\phi\phi, \text{fid}}$ is a fixed lensing potential, evaluated at the fiducial parameters given in Table I. The lensing PCs, $\Theta^{(i)}$, can be thought of as a more incisive generalization of the standard approach where a scaling parameter $C_\ell^{\phi\phi} \rightarrow A_L C_\ell^{\phi\phi}$ is added to test consistency of a model with lensing; PCs parametrize the lensing information more completely. For reasons detailed in Appendix B, the 5 best measured PCs suffice for the data and models considered here; their weights $K_\ell^{(i)}$ are shown in Fig. 11.*⁷

*⁷ Note that in Fig. 6 of [18], $K_\ell^{(i)}$ were scaled for display purposes.




FIG. 12. Joint posterior distribution for PCs $\Theta^{(1-5)}$ (black) as a product of the individual posteriors from 50 independent all-sky simulations. The joint posterior is unbiased to a small fraction of the width of the distribution of a single simulation (blue) and its Fisher prediction (red dashed).

Moreover, the PCs decouple the information on the lensing power spectrum from the parameters that control the unlensed spectrum, whereas A_L multiplies a $C_\ell^{\phi\phi}$ that depends on such parameters. PCs can therefore be more directly compared with other measurements of $C_\ell^{\phi\phi}$, most notably from lensing reconstruction using the higher point information in the temperature and polarization fields themselves.

For example, in a cosmology where the unlensed power spectra fluctuate low but the lensing potential is consistent with the underlying model, there will be a tendency for high A_L : the unlensed CMB will drive the amplitude of fluctuations A_s down, which will at the same time lower the amount of lensing predicted. To match the amount of lensing present in the data, A_L or similar lensing parameter is then increased. Hence measuring a high A_L does not necessarily signal a deviation in $C_\ell^{\phi\phi}$ itself.

In the PC approach, the unlensed CMB is described by separate parameters $\tilde{\theta}_A$. For models that modify only the low redshift physics involved in the growth of structure and cosmic acceleration it is sufficient to take these to be the equivalent of the six Λ CDM parameters. These $\tilde{\theta}_A$ change the unlensed power spectra in exactly the manner of the Λ CDM parameters, but unlike those, they have no effect on $C_\ell^{\phi\phi}$.

To summarize, CMB power spectra C_ℓ^{XY} can be effectively parameterized in terms of eleven parameters $\tilde{\theta}_A$ and $\Theta^{(i)}$ in any model that deviates from Λ CDM only after recombination, with $i = 1 - 5$ sufficing for the data and models we study. In this setup, $\Theta^{(i)}$ represent direct, optimally weighted, measurements of the lensing potential.




FIG. 13. Correlation matrix for $\Theta^{(i)}$ averaged over 50 MCMC analyses. Black squares represent ones on the diagonal. The tilded parameters affect only the unlensed CMB, as explained in the text.

B. MCMC analysis of lens and unlensed parameters

We run MCMC analyses on 50 independent lensed CMB sky simulations to check our likelihood model and to determine properties of the parameters $\tilde{\theta}_A, \Theta^{(i)}$.

So far we have made numerous assumptions, for example that our models for the likelihood and non-Gaussian covariance are correct, that the Fisher-based construction of the PCs suffices, that neglecting higher $\Theta^{(i)}$ does not affect the constraints and that agreement between theoretical and simulated power spectra is sufficient (see Appendix A). It is thus a nontrivial check of our analysis to ascertain that the constraints on $\Theta^{(i)}, \tilde{\theta}_A$ are unbiased with respect to the fiducial model. To check this, we multiplied 50 MCMC posterior probabilities for $\Theta^{(i)}, \tilde{\theta}_A$. The results for $\Theta^{(i)}$ are shown in Fig. 12 and show no significant bias relative to the standard deviation of a single MCMC posterior; the same conclusion is valid also for $\tilde{\theta}_A$.

The Fisher analysis also predicts that $\Theta^{(i)}$ as determined by the data should be uncorrelated. This assertion can be checked by averaging covariance matrix of the cosmological parameters $\Theta^{(i)}, \tilde{\theta}_A$ over the MCMC analyses. Correlation coefficients of $\Theta^{(i)}$ obtained from this covariance matrix are shown in Fig. 13. As expected, the lensing principal components are only very weakly mutually correlated. Additionally, they are also only mildly correlated with the unlensed parameters $\tilde{\theta}_A$. Most significant of these are $R = 0.18$ correlation between $\Theta^{(3)}$ and θ_* and $R = -0.14$ correlation between $\Theta^{(4)}$ and θ_* . This is somewhat counterintuitive, as θ_* shifts the angular scale of acoustic features whereas lensing mainly smears the peaks by superimposing magnified and demagnified regions. While largely true, the effect of lenses that are on scales smaller than the acoustic scale $\ell \gtrsim 200$ is not purely a smearing effect, causing a component that is slightly out of phase with the peaks in the unlensed power spectra, leading to the observed correlation.

The full results for all pairs of the 11 parameters in a single simulation are shown in Fig. 14. Note that posterior distribution for the parameters $\Theta^{(i)}, \tilde{\theta}_A$ seems to




FIG. 14. MCMC constraints on the lens and unlensed parameters $\Theta^{(i)}, \tilde{\theta}_A$ in a typical simulation with a Gaussian (black curves) and non-Gaussian covariance (red shaded) analysis.

be very well approximated by a multivariate normal distribution. This implies that the Fisher approximation should be quite accurate in this space as we explicitly verify in Fig. 12. We exploit the multivariate normal nature of the posterior in the $\Theta^{(i)}, \tilde{\theta}_A$ variables in the following sections.

C. Impact of lensing-induced covariance revisited

In Fig. 14 we also show the effect of the non-Gaussian covariance in the 11D lens and unlensed parameter space; only $\Theta^{(1)}$ and $\Theta^{(2)}$ show significant effects of neglecting the non-Gaussian likelihood. These two measurements are strongly affected, because significant portion of the noise in these measurements arises from the sample vari-

ance of the lenses. Neglecting lensing-induced terms in the data covariance effectively omits this noise contribution, which leads to overly optimistic estimates on $\Theta^{(1)}, \Theta^{(2)}$. Measurements of other lensing principal components and $\tilde{\theta}_A$ are limited by other sources of noise – instrumental noise and cosmic variance of the unlensed CMB – and the resulting constraints are thus not strongly affected by the non-Gaussian part of the covariance.

If the eleven parameters $\Theta^{(i)}, \tilde{\theta}_A$ contain all information about a particular cosmological model, a Karhunen-Loève analysis applied to the Fisher information matrices [19] reveals that non-Gaussian covariances can degrade the standard deviation of any linear combination of these parameters by maximally 2.53.*⁸ This generalizes the discussion of the most degraded linear combination M of cosmological parameters from §III. This PC based quantification of the effect of non-Gaussian covariances is not restricted to the models investigated here and can be applied to more general extensions of Λ CDM. It is also not necessary to assume a *linear* relationship between these effective parameters and the bare cosmological parameters, or the validity of the Fisher approximation for the latter.

Moreover, we can use $\Theta^{(1,2)}$ to directly translate the effect of lensing-induced covariance on cosmological parameter constraints. Those combinations of cosmological parameters which are limited by our knowledge of $\Theta^{(1,2)}$, in other words those constrained by the (mostly low ℓ) lensing information, will be strongly affected if we neglect the non-Gaussian covariance.

For example in Λ CDM, lensing information helps mainly ω_c and A_s constraints. Increasing either of these parameters increases $C_\ell^{\phi\phi}$, lensing information thus helps constrain the direction of simultaneously increasing ω_c and A_s . As we can see in Fig. 3, adding non-Gaussian covariance weakens exactly this parameter combination the most. The reason becomes clear when we examine how change in the parameter combination M , the most sensitive to the non-Gaussian covariances, projects onto the changes in the effective parameters $\Delta\Theta^{(i)}, \Delta\tilde{\theta}_A$. As expected, the main effect is a shift in $\Theta^{(1)}, \Theta^{(2)}$, which are the variables showing most of the effect of the non-Gaussian covariance; this shift is captured in Fig. 15. In comparison, shifts in the other effective parameters $\Theta^{(i)}, \tilde{\theta}_A$ are at least a factor of few smaller, as measured by the sizes of the marginalized posterior. This means that already within Λ CDM it is possible to construct a parameter combination which has a dominant effect of changing the lensing potential ($\Theta^{(1)}, \Theta^{(2)}$) and keeps the unlensed power spectra ($\tilde{\theta}_A$) relatively intact. Because of that, the relative change in the standard deviation for M brought about by the non-Gaussian covariance 2.03 is already close to the maximal possible value of 2.53.




FIG. 15. Lines show how $\Theta^{(1)}$ and $\Theta^{(2)}$ change when we increase M (black dashed), M^ν (green) or M^w (black solid); length of the lines is arbitrary. For comparison, in the background we show typical constraints on these two parameters in an MCMC analysis based on non-Gaussian (red) and Gaussian (blue) covariance.

By extending the cosmological model to Λ CDM+ $\sum m_\nu$, we increase the parameter freedom, which enables us to find a parameter combination which is slightly more limited by the lens sample covariance and shows degradation of 2.05. From the perspective of the effective parameters $\Theta^{(i)}, \tilde{\theta}_A$ this occurs because it is possible to achieve the same change of the lensing potential, as seen in the nearly identical directions of M and M^ν in Fig. 15, with a smaller change in the unlensed power spectra. This increases relative importance of the low ℓ lensing information $\Theta^{(1)}, \Theta^{(2)}$ in constraining M^ν , which directly leads to a larger impact of non-Gaussian covariance.

Moreover, lensing information is now important for three parameters $\{\omega_c, A_s, \sum m_\nu\}$ unlike two in Λ CDM. Because the degeneracy structure involves three parameters, the non-Gaussian effect is hidden from the covariance of any two, once the third is marginalized. In a three-dimensional likelihood for these three parameters, the effect of non-Gaussian covariance is clearly visible.

For Λ CDM+ w the analysis is similar. Again, by adding a new parameter on top of Λ CDM we can find M^w which shows degradation larger than what is seen in Λ CDM, in this case 2.28. For Λ CDM+ w this happens, because the projection of M^w onto $\{\Theta^{(1)}, \Theta^{(2)}\}$ is more aligned with the direction maximally impacted by the non-Gaussian covariance, see Fig. 15.

D. Consistency tests

The simplest lensing consistency test of a model is to compare lensing potential measured through $\Theta^{(i)}$ against

*⁸ Using covariance matrices of $\Theta^{(i)}, \tilde{\theta}_A$ from 10 MCMC simulations we checked this prediction is on average correct to ± 0.04 .

what is expected based on the cosmological model determined by constraints on the unlensed power spectra θ_A . This test mainly checks the internal consistency of the model assumptions.

If there are also additional measurements of the lensing potential, then a sharper consistency check is possible [18]. A certain linear combination of the power spectrum PCs,

$$\Psi^{(1)} = \sum_{i=1}^5 \mathcal{T}_i \Theta^{(i)} \quad (23)$$

with $\mathcal{T}_i = \{32.3, -15.9, 0.330, 1.72, -0.608\}$, is predicted by the Fisher analysis to be limited mainly by lens sample variance, which drops out when comparing with other measurements on the same patch of sky leaving a nearly noise-free consistency test. This consistency test therefore checks for systematics in the data analysis, foregrounds, and assumptions about the unlensed power spectra.

In Fig. 16, we illustrate this idea explicitly by comparing posterior mean values of $\Psi^{(1)}$ determined from 50 simulated power spectra using MCMC against values of $\Psi^{(1)}$ determined directly from a known realization of the lensing potential, which can be thought of as a limiting case of lensing potential measurement with no instrumental noise. The latter approach estimates $\Theta^{(i)}$ from simulated $C_\ell^{\phi\phi}$ using an unbiased estimator

$$\Theta_{\text{est}}^{(i)} = \sum_{\ell} K_{\ell}^{(i)} (\ln C_{\ell}^{\phi\phi} - \langle \ln C_{\ell}^{\phi\phi} \rangle), \quad (24)$$

where $\langle \cdot \rangle$ is expectation value over realizations, and combines the results according to (23). The observed correlation is indeed very tight; residual scatter in Fig. 16 is caused by instrumental noise and variance of the unlensed CMB which affect the power spectrum measurement.

Notice also that this direction is mainly composed of $\Theta^{(1)}$ and $\Theta^{(2)}$ in nearly the same combination that is maximally affected by non-Gaussian lens sample variance. While being limited by the lens sample variance is detrimental to cosmological parameter constraints, for consistency tests this is advantageous because lens sample variance drops out when comparing measurements on the same patch of sky.

There is also another combination of PCs, $\Psi^{(2)}$,

$$\Psi^{(2)} = \sum_{i=1}^5 \mathcal{U}_i \Theta^{(i)} \quad (25)$$

with $\mathcal{U}_i = \{31.3, 17.6, 10.4, 0.945, -1.10\}$, which can serve as a similar consistency check; this consistency check is slightly weaker than the test using $\Psi^{(1)}$ due to the larger impact of noise and sample variance of the unlensed CMB. This weakening can be seen in Fig. 17, where we compare values of $\Psi^{(2)}$ determined from 50 simulated power spectra against values from known realization of the lensing potential.




FIG. 16. Consistency mode $\Psi^{(1)}$ as determined from posterior mean values in 50 simulated lensed power spectra through MCMC analysis against values determined from known realizations of the lensing potential (see text for details). The dashed line represents points where the two determinations are equal.




FIG. 17. Same as Fig. 16 but for the consistency mode $\Psi^{(2)}$.

Notice also that there is no indication of bias in the power spectrum estimate of the two consistency modes in Figs. 16, 17. To better quantify this, in Figure 18 we show product of 50 posteriors for measurements of $\Psi^{(1)}$, $\Psi^{(2)}$ as determined from our MCMC simulations and the determination is indeed unbiased; from the power spectra side there does not seem to be any problem for the consistency check.

We also find that $\Psi^{(1)}$ is almost uncorrelated with the Λ CDM parameters describing the unlensed power spectra; the largest correlation coefficient we find is $R = -0.03$ and occurs between $\Psi^{(1)}$ and $\widetilde{\Omega_b h^2}$. The other consistency mode is slightly more correlated with the unlensed parameters; the most correlated with $\Psi^{(2)}$




FIG. 18. Joint posterior of the consistency parameters $\Psi^{(1)}$, $\Psi^{(2)}$ (black) as the product of 50 individual posterior distributions from independent all-sky simulations. Compared against the width of a single posterior (blue) there is no indication of bias with respect to the fiducial value $\Psi^{(i)} = 0$ at a fraction of the standard deviation.

are $\tilde{\theta}_*$ with correlation coefficient $R = 0.12$ and $\widetilde{\Omega_b h^2}$ with $R = 0.08$.

Failure of the consistency check could indicate an unlensed spectrum that is not described by the flat Λ CDM parameters. In fact in [18] it was shown that the second consistency mode $\Psi^{(2)}$ is correlated with the spatial curvature given their similar effects on the acoustic peaks. We run a single MCMC analysis in which we added $\tilde{\Omega}_K$ to the unlensed parameters and confirm this finding; there is a strong correlation between $\tilde{\Omega}_K$ and $\Psi^{(2)}$ with correlation coefficient $R = -0.62$. In a non-flat Universe analyzed as flat this would lead to failure in the consistency check as $\Psi^{(2)}$ would move from its true value to absorb the unaccounted for curvature. The main consistency mode, $\Psi^{(1)}$, is correlated with $\tilde{\Omega}_K$ at the $R = 0.29$ level and is a weaker check on curvature.

E. Effective Likelihood for Model Building

Given the nearly multivariate normal posterior probability of the effective parameters $d_a = \{\theta_A, \Theta^{(i)}\}$, we can also use a single MCMC analysis to compress the whole CMB power spectra data into 11 numbers for the mean and the 11×11 covariance matrix of d_a . These can be used to form an effective likelihood function $\mathcal{L}_{\text{eff}}(\bar{d}_a | \theta_A)$ defined as

$$-2 \ln \mathcal{L}_{\text{eff}} = \sum_{ab} [\bar{d} - d(\theta_A)]_a (\text{Cov}_d^{-1})_{ab} [\bar{d} - d(\theta_A)]_b. \quad (26)$$

Here θ_A models the expectation values for the data \bar{d}_A . This effective likelihood can be now used to probe a broad class of models without any explicit use of the raw CMB

power spectra data. This class contains not just the extensions to Λ CDM considered here but also models which are indistinguishable from Λ CDM at recombination and for which CMB lensing is the dominant source of information on the extension. For example, many models of dark energy and modified gravity fall into this class, if we are willing to ignore the extra information coming from the integrated Sachs-Wolfe effect and other secondaries. In principle, the technique can be extended to incorporate such effects by extending the set of unlensed parameters $\tilde{\theta}_A$.

In the context of model building, one can envision a scenario where Λ CDM parameters produce a poor effective likelihood for the data and motivate the extensions. The effective likelihood can then be used as a quick spot check as to whether the given extension improves the fit.

Let us illustrate this technique on the extensions of Λ CDM considered above. First, it is necessary to find the functional dependence of d_a on the cosmological parameters θ_A . The values of the unlensed parameters $\tilde{\theta}_A$ are the same as the true cosmological parameters θ_A , while values of the lensing principal components $\Theta^{(i)}$ can be determined directly from the definition (24) given $C_\ell^{\phi\phi}$ alone. The full parameter space of the given extension can then be explored with an MCMC in the general case where $d_a(\theta_A)$ is nonlinear across the allowed region of the parameter space as in the Λ CDM+ w extension. In a case such as Λ CDM+ $\sum m_\nu$, where the mapping can be linearized, it is possible to get a good estimate of parameter constraints even without performing any additional MCMC.

In Fig. 19 we show comparison of Λ CDM+ w parameter constraints obtained in the analysis presented in the previous paragraph against results of the standard MCMC analysis. Because the mapping onto the effective parameter space is non-linear, it is necessary to perform an additional MCMC run. Note that this mapping alone accounts for most of the non-normal posterior probability in the (ω_c, w) plane despite being based on a normal distribution for the effective parameters. It slightly underestimates the lower limit on w , presumably due to the neglect of the integrated Sachs-Wolfe effect in the unlensed parameters. For Λ CDM+ $\sum m_\nu$, the agreement between the simplified and standard analyses is even better.

V. DISCUSSION

Future measurements of lensed CMB power spectra will be increasingly affected by the lens sample variance and its effect on parameter constraints will have to be included into the analysis pipelines once polarization measurements approach the sample variance limit. To this end we have developed and tested such an analysis pipeline starting from simulated lensed maps with CMB-S4 level instrument noise through to constraints on cosmological parameters and the lens power spectrum.




FIG. 19. Comparison of constraints on Λ CDM+ w parameters from the standard analysis (red shaded) with results of an approximate analysis based on the effective likelihood of $\{\Theta^{(i)}, \hat{\theta}_A\}$ instead of the raw CMB data (black).

The first piece in the pipeline is a model for the likelihood function of full-sky lensed CMB power spectrum data, which includes the non-Gaussian effects of lensing sample variance that correlate the measurements. This model considers large and small multipole data as independent. At small multipoles, the likelihood assumes data are Wishart distributed and drops the small non-Gaussian effects of lensing, while at large multipoles it assumes they are multivariate-normal distributed, with a covariance matrix which includes the lensing correlation across multipole moments.

With this likelihood we investigated parameter constraints from simulated lensed CMB maps of a fiducial Λ CDM model. We obtain MCMC parameter constraints on the Λ CDM parameters as well as two extensions, Λ CDM+ $\sum m_\nu$ and Λ CDM+ w . The dominant effect of the lensing-induced covariance in all of the models is more than four-fold increase in variance of particular combinations of cosmological parameters M, M^ν, M^w . As a consequence, if the analysis of the real data was performed using Gaussian covariance in the likelihood, instead of the proper non-Gaussian covariance, there is a high chance of committing Type 1 error - mistakenly ruling out true cosmological model. This would potentially affect concordance studies comparing constraints from various datasets. Shifts in the best fit basis parameters and change in constraints of the other parameter combinations are typically not as significant due to marginalization. The exception is Λ CDM+ w where a significant

degradation in the lower limit for w is manifest in the MCMC results. This degradation is hidden from the local Fisher forecasts as well as previous studies due to the strongly non-normal posterior distribution of w .

Then we explored the use of direct constraints on the lensing potential through a principal component analysis. Here eleven parameters effectively describe most of the cosmological information contained in the lensed CMB power spectra. Five of these parameters, $\Theta^{(i)}$, are the best measured principal components of the lensing potential while the remaining six, $\hat{\theta}_A$, parameterize the unlensed power spectra. Measurement of $\Theta^{(i)}$ from data are well suited for various consistency tests involving measurements of the lens power spectrum. This should be contrasted with the standard approach where cosmological parameters are augmented by a scaling parameter A_L to the lens power spectrum but the latter itself depends on cosmological parameters which are subsequently marginalized. Here the measured $\Theta^{(i)}$ can be compared directly against lensing potentials corresponding to the measured unlensed parameters $\hat{\theta}_A$, to check the internal consistency of a particular cosmological model

On 50 MCMC analyses we tested our PC analysis pipeline, not finding any significant bias in either $\Theta^{(i)}$ or $\hat{\theta}_A$. The lensing principal components $\Theta^{(i)}$ are found to be only weakly correlated, both mutually and with the unlensed parameters $\hat{\theta}_A$. The majority of the effects of the non-Gaussian covariance consists of degrading constraints on the two leading lensing principal components, $\Theta^{(1)}$ and $\Theta^{(2)}$ (see Fig. 14). This allows explanation of the parameter constraint degradations seen in the cosmological models - in each there is a parameter combination which is predominantly limited by the low ℓ lensing information which has large lens sample variance relative to other sources of noise. Neglecting non-Gaussian terms in the covariance effectively neglects this source of noise, which misleads the parameter constraints.

The effect of lens sample variance on the PCs enables a sharp consistency test against other measurements of the lens power spectrum. The degradation in parameter errors reflects a linear combination of PCs whose measurements are nearly lens sample variance limited. Independent measurements on the same patch of sky, e.g. through direct reconstruction from the CMB four point functions, should agree since the sample variance is common to both. In this paper we checked that the combinations of $\Theta^{(i)}$ which are expected to form the most stringent consistency tests, $\Psi^{(1)}$ and $\Psi^{(2)}$, satisfy theoretical expectations. Their values determined from lensed power spectra using MCMC analyses are correlated with the “true” values determined from the known realization of $C_\ell^{\phi\phi}$. They are also unbiased and nearly uncorrelated with the unlensed parameters $\hat{\theta}_A$.

Failure of any of these consistency tests may indicate new physics beyond flat Λ CDM, as we show on an example of spatial curvature. Constraints on $\Theta^{(i)}, \hat{\theta}_A$ can also be used for model building purposes given their simple multivariate normal form. As illustrated using the

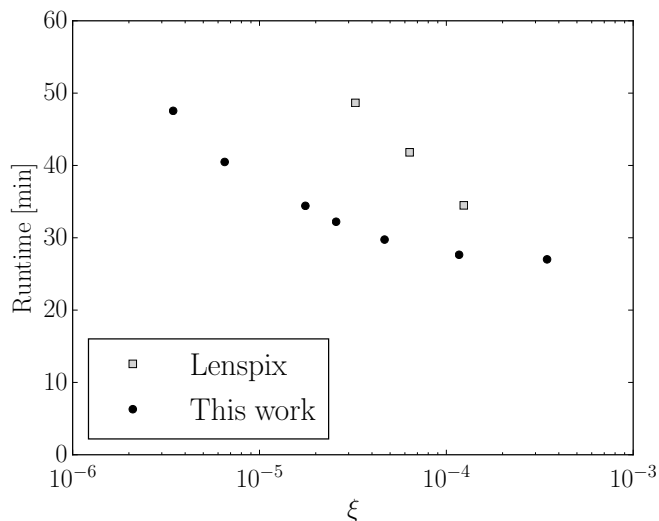


FIG. 20. Bias ξ caused by the interpolation part of the lensing algorithm and corresponding runtime for various values of the precision parameter `interp_factor`. Comparison of the Lenspix interpolation routine (gray squares, from right `interp_factor` values 2, 2.5 and 3) and our modifications (black dots, values 2, 2.5, 3, 3.5, 4, 5 and 6).

dark energy equation of state, one can rapidly explore lensing constraints on extensions to Λ CDM using an effective likelihood for these parameters without recourse to the original CMB power spectrum data or experimental specifics.

ACKNOWLEDGMENTS

We thank Chen He Heinrich, Alessandro Manzotti and Marco Raveri for useful discussions. This work was supported by NASA ATP NNX15AK22G and the Kavli Institute for Cosmological Physics at the University of Chicago through grant NSF PHY-1125897 and an endowment from the Kavli Foundation and its founder Fred Kavli. WH was additionally supported by U.S. Dept. of Energy contract DE-FG02-13ER41958 and the Simons Foundation. PM acknowledges the hospitality of Nordita where part of this work was completed during its program Advances in Theoretical Cosmology in Light of Data. We acknowledge use of the CAMB, Lenspix and CosmoMC software packages. This work was completed in part with resources provided by the University of Chicago Research Computing Center.

Appendix A: Simulated CMB sky and Lenspix modifications

To simulate lensed CMB data we modify publicly available code Lenspix^{*9} [25]. In this code the unlensed CMB is first evaluated on a high resolution equicylindrical grid. The lensed CMB is then evaluated on a lower resolution Healpix grid [28] through a remapping by a deflection field, determined by a gradient of the lensing potential ϕ . Values of the unlensed CMB at points which are remapped onto the Healpix grid points are obtained using a bi-cubic interpolation from the high resolution grid. Our simulations are run with precision parameters `nside` = 4096 and `lmax` = 8000.

The precision with which the code calculates the lensed power spectra depends on the point density of the high resolution grid, which is parameterized by an oversampling factor `interp_factor`. Simulations with `interp_factor` ~ 2 , which were the largest we could originally run on a single node of our computer cluster due to limited memory, lead to lensed CMB power spectra biased at high ℓ . Such bias leads to ~ 0.2 standard deviations shift in the likelihood function in the $\ln A_s$ direction (with other cosmological parameters fixed); we did not investigate other parameters in depth but in general parameters constrained by high ℓ data are sensitive to this bias.

This power spectra bias can be quantified by a parameter

$$\xi = \frac{C_{3000}^{TT} \Big|_{\phi=0} - \tilde{C}_{3000}^{TT}}{\tilde{C}_{3000}^{TT}}, \quad (\text{A1})$$

relative difference at $\ell = 3000$ of the temperature power spectrum C^{TT} “lensed” in Lenspix by a zero-deflection field and the unlensed temperature power spectrum \tilde{C}^{TT} . If interpolation was exact ξ would vanish. However, the unlensed and lensed CMB are evaluated at different grids and interpolation leads to numerical bias even when there is no lensing present. This bias appears to be – up to cosmic variance – independent of the cosmology and comparable in temperature and polarization. It typically grows with investigated multipole, for comparison we therefore choose the largest data multipole considered in the paper, $\ell = 3000$.

To overcome the large ξ bias and avoid the related shifts in the likelihood function, we modify the code such that it works only with smaller portions of the high resolution map of unlensed CMB at any given time and never stores the whole map in memory. This allows us to run with higher values of `interp_factor` and achieve smaller values of ξ .

We further replace the original high precision calculation of partial derivatives of the unlensed CMB variables,

^{*9} <https://github.com/cmbant/lenspix>




FIG. 21. Lensed power spectra bias b_ℓ^{XY} for several values of ℓ , averaged over 2000 lensed CMB simulations calculated with the precision settings used in this work (black). In red the same quantities determined from 400 lensed CMB simulations calculated with original Lenspix interpolation algorithm with `interp_factor = 2`. Error bars represent errors on the mean estimated from the simulated values.

which is part of the Lenspix interpolation algorithm, by a less precise (for a given high resolution grid) but significantly faster routine. This enables us to obtain higher interpolation precision without sacrificing runtime by increasing the density of the high resolution grid of the unlensed CMB.

Finally, the precision of variables describing the angular positions of the points in the high resolution grid is increased to avoid certain artifacts in lensed CMB maps.

We compare values of ξ and runtime for several values of `interp_factor` with the original and simplified calculation of the partial derivatives in Fig. 20. It is clear that although the original routine is superior for a fixed high resolution grid, for a fixed runtime it is advantageous to use a simpler partial derivative calculation and increase the density of the grid. Simulations used in this work were calculated with `interp_factor = 4`.

To judge agreement between the lensed power spectra from simulations and the theoretical expectation calculated by CAMB, we define bias variables

$$b_\ell^{XY} = \frac{\sum_{\ell'} \ell'(\ell' + 1) \Delta C_{\ell'}^{XY}}{\sum_{\ell'} \ell'(\ell' + 1) C_{\ell'}^{XY, \text{fid}}}, \quad (\text{A2})$$

which can be evaluated for each simulated CMB sky. Here the sums go over a bin of width $\Delta\ell = 200$ centered on ℓ and $\Delta C_{\ell'}^{XY}$ is a difference between simulated and expected value of power spectra, defined in (9). In Fig. 21 we plot average values of b_ℓ^{XY} from simulations for several values of ℓ ; we show the levels of bias achieved with both the simulations settings used in this work and the original Lenspix code with `interp_fact = 2`. In the latter, a small bias is visible for $XY = TT, EE$.

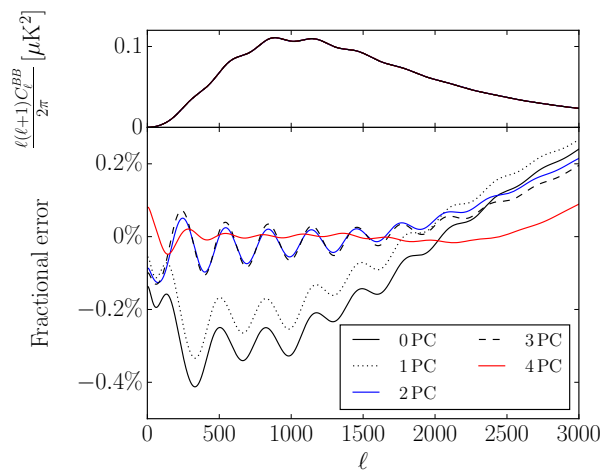


FIG. 22. Lensed C_ℓ^{BB} calculated with lensing potential increased by $\Delta C_\ell^{\phi\phi}$ which corresponds to parameter shifts listed in Tab. II and its representation in terms of the first $N \in 0 \dots 4$ PCs (top: absolute; bottom percent error between the two).




FIG. 23. Dependence of $\chi_{\text{PC},N}^2$ (B2), measure of error caused by approximating the lensing potential using first N principal component, on N . In this case $\Delta C_\ell^{\phi\phi}$ corresponds to parameter shifts listed in Tab. II.

Appendix B: Number of lensing principal components

Two considerations guide the choice of the number of principal components of the lensing potential to be measured from the lensed CMB power spectra. Keeping a larger number of PCs leads to a more accurate description of the lensed power spectra. On the other hand, increasing number of parameters slows down convergence of the MCMC calculations, requiring physicality priors, since the higher PCs are more poorly constrained by definition. In this section we justify our choice of using five PCs.

TABLE II. Shifts in the cosmological parameters used to probe approximations of the lensing potential in terms of lensing PCs

| Parameter | Shift |
|----------------|------------------------|
| h | 0.175 |
| $\Omega_c h^2$ | -1.0×10^{-3} |
| $\Omega_b h^2$ | 3.5×10^{-5} |
| n_s | 1.2×10^{-3} |
| A_s | -1.6×10^{-11} |
| τ | -2.2×10^{-3} |
| w | -0.52 |

First we look at the fidelity in reproducing lensing effects in the observed C_ℓ^{XY} power spectra. For definiteness, we perturb the fiducial $C_\ell^{\phi\phi}$ by a $\Delta C_\ell^{\phi\phi}$ which corresponds to shifts in the cosmological parameters given in Table II (at fixed unlensed power spectra). This change represents a realistic change in the lensing potential which might be encountered in a real analysis, as cosmological model with parameters from Tab. II is between 68% and 95% probability contours for Λ CDM+ w model in the simulation investigated in the main text. For completeness and comparison to the results in the main text, the lensing principal components arising from this change are

$$\Theta^{(i)} = \{-0.012, 0.11, -0.011, -0.06, -0.012\}. \quad (\text{B1})$$

Note that the investigated parameter change is not aligned with change of M^w from (B1), nor need it be since it represents the degeneracy direction rather than the direction most affected by lens covariance.

In top panel of Fig. 22 we show the resulting C_ℓ^{BB} power spectrum calculated with lensing potential

changed by the full $\Delta C_\ell^{\phi\phi}$ vs. when this change is approximated using the first $N \in \{0, 1, 2, 3, 4\}$ lensing PCs; the difference is too small to be visible directly and so the bottom panel shows the percent error. Note that the smallness of these changes explains why in the main text approximations based on linearizing power spectra deviations and PC amplitudes are excellent even for relatively large cosmological parameter shifts.

Since the change in Eq. B1 is dominated by $\Theta^{(2)}$, most of the improvement in fidelity comes when adding that component. In fact for this particular $\Delta C_\ell^{\phi\phi}$, first two PCs are sufficient to faithfully describe effects of lensing in TT , EE and TE power spectra extremely well. The next large jump in fidelity comes with the fourth PC which is associated with the high multipole range of BB in Fig. 22. We checked several other choices of allowed $\Delta C_\ell^{\phi\phi}$ and for all of them four principal components lead to small errors on the power spectra level.

To quantify the total significance of the errors we construct

$$\chi_{\text{PC},N}^2 = \sum_{i,j} \delta D_i(N) (\text{Cov}_{i,j})^{-1} \delta D_j(N), \quad (\text{B2})$$

where for brevity we introduce $\delta D_i(N) = \delta C_\ell^{XY}(N)$, the power spectrum error caused by approximating $\Delta C_\ell^{\phi\phi}$ using the first N PCs, with i indexing all multipoles and power spectra types. As N increases, the PCs approximate the full effect of lensing better and $\chi_{\text{PC},N}^2$ decreases. In Fig. 23 we show this dependence; as we saw before, adding fourth PC leads to a significant improvement in our ability to capture the effects of gravitational lensing on the CMB. For some choices of $\Delta C_\ell^{\phi\phi}$, adding fifth PC improves $\chi_{\text{PC},N}^2$ by a factor of a few on top of the \sim hundred-fold improvement in $\chi_{\text{PC},N}^2$ arising from using four PCs. For this work we decided to include fifth PC into the analysis as well, even though its inclusion is not expected to have any significant impact on the results.

-
- [1] P. A. R. Ade *et al.* (Planck), *Astron. Astrophys.* **594**, A13 (2016), arXiv:1502.01589 [astro-ph.CO].
- [2] W. Hu, *Phys. Rev.* **D64**, 083005 (2001), arXiv:astro-ph/0105117 [astro-ph].
- [3] K. M. Smith, O. Zahn, and O. Dore, *Phys. Rev.* **D76**, 043510 (2007), arXiv:0705.3980 [astro-ph].
- [4] A. Lewis and A. Challinor, *Phys. Rept.* **429**, 1 (2006), arXiv:astro-ph/0601594 [astro-ph].
- [5] C. M. Hirata, S. Ho, N. Padmanabhan, U. Seljak, and N. A. Bahcall, *Phys. Rev.* **D78**, 043520 (2008), arXiv:0801.0644 [astro-ph].
- [6] D. Hanson *et al.* (SPTpol), *Phys. Rev. Lett.* **111**, 141301 (2013), arXiv:1307.5830 [astro-ph.CO].
- [7] S. Das *et al.*, *Phys. Rev. Lett.* **107**, 021301 (2011), arXiv:1103.2124 [astro-ph.CO].
- [8] R. Keisler *et al.*, *Astrophys. J.* **743**, 28 (2011), arXiv:1105.3182 [astro-ph.CO].
- [9] P. A. R. Ade *et al.* (Planck), *Astron. Astrophys.* **571**, A17 (2014), arXiv:1303.5077 [astro-ph.CO].
- [10] R. Keisler *et al.* (SPT), *Astrophys. J.* **807**, 151 (2015), arXiv:1503.02315 [astro-ph.CO].
- [11] P. A. R. Ade *et al.* (Planck), *Astron. Astrophys.* **594**, A15 (2016), arXiv:1502.01591 [astro-ph.CO].
- [12] P. A. R. Ade *et al.* (BICEP2, Keck Array), (2016), arXiv:1606.01968 [astro-ph.CO].
- [13] B. D. Sherwin *et al.*, *Phys. Rev.* **D95**, 123529 (2017), arXiv:1611.09753 [astro-ph.CO].
- [14] B. A. Benson *et al.* (SPT-3G), *Proceedings, SPIE Astronomical Telescopes + Instrumentation 2014: Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII: Montreal, Quebec, Canada, June 24-27, 2014*, Proc. SPIE Int. Soc. Opt. Eng. **9153**, 91531P (2014), arXiv:1407.2973 [astro-ph.IM].
- [15] S. W. Henderson *et al.*, *Proceedings, 16th International Workshop on Low Temperature Detectors (LTD 16): Grenoble, France, July 20-24, 2015*, *J. Low. Temp. Phys.* **184**, 772 (2016), arXiv:1510.02809 [astro-ph.IM].

- [16] W. L. K. Wu *et al.*, *Proceedings, 16th International Workshop on Low Temperature Detectors (LTD 16): Grenoble, France, July 20-24, 2015*, *J. Low. Temp. Phys.* **184**, 765 (2016), arXiv:1601.00125 [astro-ph.IM].
- [17] K. N. Abazajian *et al.* (CMB-S4), (2016), arXiv:1610.02743 [astro-ph.CO].
- [18] P. Motloch, W. Hu, and A. Benoit-Lévy, *Phys. Rev.* **D95**, 043518 (2017), arXiv:1612.05637 [astro-ph.CO].
- [19] K. M. Smith, W. Hu, and M. Kaplinghat, *Phys. Rev.* **D74**, 123002 (2006), arXiv:astro-ph/0607315 [astro-ph].
- [20] A. Benoit-Lévy, K. M. Smith, and W. Hu, *Phys. Rev.* **D86**, 123008 (2012), arXiv:1205.0474 [astro-ph.CO].
- [21] D. Green, J. Meyers, and A. van Engelen, (2016), arXiv:1609.08143 [astro-ph.CO].
- [22] J. Peloton, M. Schmittfull, A. Lewis, J. Carron, and O. Zahn, (2016), arXiv:1611.01446 [astro-ph.CO].
- [23] L. Knox, *Phys. Rev. D* **52**, 4307 (1995), astro-ph/9504054.
- [24] A. Lewis, A. Challinor, and A. Lasenby, *Astrophys. J.* **538**, 473 (2000), arXiv:astro-ph/9911177 [astro-ph].
- [25] A. Lewis, *Phys. Rev.* **D71**, 083008 (2005), arXiv:astro-ph/0502469 [astro-ph].
- [26] A. Lewis and S. Bridle, *Phys. Rev.* **D66**, 103511 (2002), arXiv:astro-ph/0205436 [astro-ph].
- [27] A. Gelman and D. B. Rubin, *Statist. Sci.* **7**, 457 (1992).
- [28] K. M. Gorski, B. D. Wandelt, F. K. Hansen, E. Hivon, and A. J. Banday, (1999), arXiv:astro-ph/9905275 [astro-ph].