Effect of non-Gaussian lensing deflections on CMB lensing measurements

Vanessa Böhm,[1,2] Blake D. Sherwin,[3] Jia Liu,[4] J. Colin Hill,[5,6] Marcel Schmittfull,[6] and Toshiya Namikawa[7]

[1]*Berkeley Center for Cosmological Physics, University of California, Berkeley, California 94720, USA*
[2]*Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, California 93720, USA*
[3]*Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Wilberforce Road, Cambridge CB3 0WA, United Kingdom*
[4]*Department of Astrophysical Sciences, Princeton University, Princeton, New Jersey 08544, USA*
[5]*Center for Computational Astrophysics, Flatiron Institute, 162 5th Avenue, New York, New York 10010, USA*
[6]*Institute for Advanced Study, Einstein Drive, Princeton, New Jersey 08540, USA*
[7]*Leung Center for Cosmology and Particle Astrophysics, National Taiwan University, Taipei 10617, Taiwan*

(Received 6 June 2018; published 13 December 2018)

We investigate the impact of non-Gaussian lensing deflections on measurements of the CMB lensing power spectrum. We find that the false assumption of their Gaussianity significantly biases these measurements in current and future experiments at the percent level. The bias is detected by comparing CMB lensing reconstructions from simulated CMB data lensed with Gaussian deflection fields to reconstructions from simulations lensed with fully non-Gaussian deflection fields. The non-Gaussian deflections are produced by ray tracing through snapshots of an N-body simulation and capture both the non-Gaussianity induced by nonlinear structure formation and by multiple correlated deflections. We find that the amplitude of the measured bias can be modeled with analytic expressions for a lensing bispectrum-induced bias derived by Böhm *et al.* in 2016 when post-Born corrections are included in the lensing bispectrum model. The bias is largest in temperature-based measurements, where it is detected with a significance of 2.84σ in the power spectrum of reconstructed convergence fields. Cross-correlating the reconstruction with the noiseless input convergence fields results in a 5.21σ detection. We do not find evidence for the bias in measurements from a combination of polarization fields (EB, EB). We argue that this non-Gaussian bias should be even more important for measurements of cross-correlations of CMB lensing with low-redshift tracers of large-scale structure.

DOI: 10.1103/PhysRevD.98.123510

I. INTRODUCTION

Photons of the cosmic microwave background (CMB) get deflected by the cosmic matter distribution between the surface of last scattering and the observer. This effect is known as CMB lensing (see e.g., Refs. [1,2] for reviews). Coherent deflections distort the observed CMB fluctuations in both temperature and polarization in a characteristic way. The statistics of the deflections contain a vast amount of cosmological information. They are sensitive to cosmological parameters that determine the formation of cosmic structure, such as a combination of σ_8 and Ω_m, the sum of neutrino masses [3] and the presence of dark energy [4]. They are also a probe of the flatness of space because curvature changes the relative efficiency of lensing events at different distances. Different to other probes of large-scale structure, CMB lensing is mostly sensitive to structures at relatively high redshifts ($z \approx 2$) and has the advantage of directly probing the total matter distribution.

Since the first detection of CMB lensing in cross-correlations [5,6], CMB lensing measurements have matured from detections in CMB data alone [7], through increasingly significant detections in CMB temperature, polarization and cross-correlations [8–12] to a compatible and complementary cosmological probe [13–15]. Forecasts for current and future surveys [16–19] promise sample variance limited measurements of the CMB lensing power spectrum up to multipoles of $L \approx 1000$ and a sensitivity of these measurements to the total mass of neutrinos of $\sigma_{\sum_\mathrm{m}} \approx 30$ meV if combined with suitable other probes to break degeneracies with τ and $\Omega_\mathrm{m}h^2$.

Common CMB lensing reconstruction uses a quadratic, weighted combination of CMB fields to recover the deflection field [20,21]. Power spectrum measurements from this quadratic estimator extract lensing information from the lensed CMB four-point function. The four-point estimator for the CMB lensing power spectrum is a biased estimator. It is nonzero even in the absence of lensing and carries bias terms at all orders in the lensing power spectrum [22,23]. Other sources of systematics in CMB lensing measurements are masking, anisotropic beam or

noise properties [24,25] and foregrounds [26–29]. Biases to power spectrum measurements can either be estimated and subtracted or alleviated by suitable modifications to the lensing estimator [30–33].

Recently, Ref. [34] (hereafter BSS16) has identified a new bias to CMB lensing measurements, which arises as a consequence of the non-Gaussian structure of the lensing deflection field. BSS16 specifically considered the effect of a nonvanishing bispectrum of the lensing potential. In a purely analytic study, they found that the bispectrum that arises as a consequence of nonlinear structure formation can change the amplitude of the CMB lensing power spectrum measured from CMB temperature data in current and future experiments at the percent level. As most of these experiments rely primarily on temperature, a correction to CMB lensing measurements of this magnitude would constitute a significant systematic and, if uncorrected, hinder the accurate estimation of cosmological parameters. However, the theory calculation in BSS16 made a number of nontrivial assumptions (see Appendix A for details), and the actual size of the bias depends on their validity. In particular, BSS16 modeled the lensing bispectrum in Born approximation and did not take into account multiple deflections as another source of non-Gaussianity. Recently, Ref. [35] pointed out that the bispectrum from post-Born terms is on the same order of magnitude as the bispectrum from nonlinear gravitational evolution. The sign of these two contributions is partly opposite (depending on the triangle configuration), which suggests that including post-Born corrections results in a smaller bias than predicted by BSS16.

Motivated by this, we study the effect of non-Gaussianity on CMB lensing measurement in this work in a completely independent way with ray-traced lensing simulations. Specifically, we look at the difference between lensing power spectra measured with the standard four-point estimator in two different sets of simulated noisy, lensed CMB maps: one set lensed with purely Gaussian deflection fields and the other with fully non-Gaussian deflections obtained from ray tracing through snapshots of an N-body simulation. By using the same unlensed CMB and detector noise realizations for both sets, any significant difference in the measured spectra is a consequence of the non-Gaussianity of the deflection field and can be interpreted as a non-Gaussian bias. Although the study with simulations provides less intuition about the specific source of a non-Gaussian bias, it is in some sense more complete than the theoretical analysis in BSS16 because it captures the full non-Gaussianity of the field, which can manifest itself in more ways than a nonzero bispectrum and relies on fewer simplifying assumptions.

We compare the measured non-Gaussian bias to the theoretical predictions computed from the analytic expressions derived in BSS16. Different from BSS16 we also take into account the lensing bispectrum sourced by multiple correlated deflections (so-called post-Born corrections) in the theory. We use the analytical expression derived in Ref. [35] for modeling these corrections. Theoretical results for the bias computed with this new bispectrum model can be found in Appendix A. For a CMB-S4-like experiment, they suggest a cumulative bias over noise of $\sim 3\sigma$ in lensing reconstruction from temperature only (assuming a bin width of $\Delta L = 100$). The maximum bias-over-noise per L-bin for this configuration is 1 and occurs in the lowest bin (centered on $L = 150$).

Earlier works that have studied post-Born corrections to lensing observables in simulations include Refs. [36,37]. Reference [36] focused on the effect of post-Born corrections on parameter constraints from weak galaxy lensing measurements and found significant biases to parameters inferred from the skewness and kurtosis of the convergence field. Reference [37] quantified the impact of post-Born corrections on the CMB lensing and CMB power spectra and estimated the significance with which lensing curl modes induced by multiple deflections can be measured in future experiments. Different from these works, we study the bias induced by the non-Gaussianity of the CMB lensing deflection field on measurements of its power spectrum with the standard four-point estimator. Due to the special form of this estimator, it is sensitive to higher-order lensing correlation functions and thus sensitive to the non-Gaussianity sourced by post-Born corrections. The effect of nonlinear structure formation on lensing reconstructions, in particular, its impact on the second-order lensing bias $N^{(2)}$, was measured in ray-traced simulations for the interpretation of data from the South Pole Telescope [8] but found to be irrelevant for this specific data set. Parallel to the work presented here, Ref. [38], have carried out a measurement of a non-Gaussian bias on an independent set of ray-traced lensing simulations.

This paper is organized as follows: we start with briefly reviewing CMB lensing and CMB lensing reconstruction in Sec. II. In Sec. III we give a full overview of the production of mock CMB data maps: in subsections we describe the production of ray-traced lensing maps and their Gaussian counterparts (Sec. III A), the generation of noisy, lensed CMB simulations (Sec. III B) and the reconstruction from these mock data sets (Sec. III C). Results and their comparison to theory are presented in Sec. IV. We conclude with a discussion of the results and a comment on the importance of the non-Gaussian bias for cross-correlations with low-redshift tracers in Sec. V. For details on the theoretical bias model derived in BSS16, we refer the reader to Appendix A and Ref. [34].

II. CMB LENSING AND CMB LENSING RECONSTRUCTION

Lensing distortions are a measure of the integrated mass distribution along the photons' trajectories. In a flat standard cosmology and under the Born approximation, the

lensing convergence $\kappa(\mathbf{L})$ is related to the density contrast $\delta(\mathbf{L},\chi)$ through the line-of-sight integration,

$$\kappa(\mathbf{L}) = \frac{3}{2}\frac{\Omega_m H_0^2}{c^2}\int_0^{\chi_{\rm CMB}} \mathrm{d}\chi W(\chi,\chi_{\rm CMB})\delta(\mathbf{L},\chi), \quad (1)$$

with lensing kernel,

$$W(\chi,\chi_{\rm CMB}) = [1+z(\chi)]\frac{\chi(\chi_{\rm CMB}-\chi)}{\chi_{\rm CMB}}. \quad (2)$$

Throughout this paper, we use the flat sky approximation, where $\mathbf{L}$ denotes the wave vector of a 2D Fourier mode on the sky. Relating convergence and density according to Eq. (1) assumes that the lensing kernel weighted integral along the photon geodesic over the second derivative of the gravitational potential in the radial direction is negligible compared to the same integral over second derivatives in transverse directions. This approximation is valid for small angular scales and if the typical scale of the integration kernel is much larger than the typical fluctuation scale of the integrated quantity [39,40]. For CMB lensing the characteristic scale of the kernel is roughly the Hubble distance, and the angular scales considered in this work are small ($L > 100$).

The mapping between unlensed CMB fields (T, Q, U) and their lensed counterparts $(\tilde{T}, \tilde{Q}, \tilde{U})$ is determined by the lensing deflection angle $\boldsymbol{\alpha}$,

$$\tilde{T}(\mathbf{x}) = T[\mathbf{x}+\boldsymbol{\alpha}(\mathbf{x})], \quad (3)$$

which is to good approximation curl-free and can be expressed in terms of a scalar lensing potential $\phi(\mathbf{x})$,

$$\boldsymbol{\alpha}(\mathbf{x}) = \nabla\phi(\mathbf{x}). \quad (4)$$

Similar to overdensity and gravitational potential in three dimensions, the lensing convergence [Eq. (1)] and the lensing potential are related by the Poisson equation,

$$\kappa(\mathbf{x}) = -\frac{1}{2}\nabla^2\phi(\mathbf{x}). \quad (5)$$

We use the extended Limber approximation [41] when computing expectation values in our theory expressions. Although the lensing simulation does not make use of the Limber approximation, the mismatch on scales with $L > 100$, which we consider, is expected to be negligible [42–44].

CMB lensing reconstruction is the recovery of the lensing deflection field from lensed, noisy CMB data. It is commonly performed with an estimator that is quadratic in the lensed CMB [20,21,45],

$$\hat{\kappa}(\mathbf{L}) = \frac{1}{2}L^2 A_L^{XY}\int_{\mathbf{l}} g_{\mathbf{l},\mathbf{L}}^{XY}\tilde{X}_{\rm expt}(\mathbf{l})\tilde{Y}_{\rm expt}(\mathbf{L}-\mathbf{l}). \quad (6)$$

In Eq. (6) X and Y represent either temperature (T) or polarization fields (E/B), and the subscript "expt" labels noisy, beam-deconvolved data. The weight g and the normalization A_L depend on the fiducial lensed CMB power spectra as well as the beam and noise properties of the experiment (see Ref. [21] for the exact expressions[1]). Weight and normalization are chosen such that the estimator in Eq. (6) has minimum variance and is unbiased in the absence of any source of mode coupling other than lensing.

A few alternatives to the quadratic estimator have been proposed. Some are based on maximizing the CMB lensing posterior[2] or sampling the joint distribution of lensing deflections and CMB [47–49]. Other estimators are derived from a configuration-space perspective and use the magnification and shear of the lensed CMB fluctuations to estimate the lensing field [50–52]. To date the quadratic estimator remains the most widely used and best understood estimator for the CMB lensing deflection field.

Measurements of the CMB lensing power spectrum from the quadratic estimator are sensitive to the lensed CMB four-point function,

$$\begin{aligned}&\hat{C}^{\kappa\kappa}_{WX,YZ}(L)\\ &= \frac{1}{4}L^4 A_L^{WX}A_L^{YZ}\int_{\mathbf{l}_1,\mathbf{l}_2} g_{\mathbf{l}_1,\mathbf{L}}^{WX}g_{\mathbf{l}_2,\mathbf{L}}^{YZ}\\ &\quad\times\langle\tilde{W}_{\rm expt}(\mathbf{l}_1)\tilde{X}_{\rm expt}(\mathbf{L}-\mathbf{l}_1)\tilde{Y}_{\rm expt}(-\mathbf{l}_2)\tilde{Z}_{\rm expt}(\mathbf{l}_2-\mathbf{L})\rangle, \end{aligned}\quad (7)$$

where the expectation value is computed by averaging over orientations of the wavevector $\mathbf{L}$. As the response of the CMB to lensing is nonlinear in the deflection, this four-point estimator gets contributions from terms at all orders in the lensing convergence. Only one of the contributing second-order terms gives rise to the convergence power spectrum. The remaining terms are bias terms that need to be subtracted in order to obtain an unbiased estimate for $C_L^{\kappa\kappa}$. They are commonly summarized and labeled by their power in the lensing power spectrum: $N_L^{(0)}$ for the bias that is sourced by Gaussian CMB fluctuations (this term is present even in the absence of lensing) and $N_L^{(1)}$ for all biases proportional to $C_L^{\kappa\kappa}$ and $N_L^{(2)}$ for biases proportional to $(C_L^{\kappa\kappa})^2$ [22,23]. The $N_L^{(2)}$ bias can be greatly reduced by a slight modification to the lensing weights, see e.g.,

[1]Reference [21] uses unlensed power spectra in the lensing weights. Replacing them by their lensed counterparts partly removes higher-order biases from the power spectrum estimate [23,30,46].

[2]The quadratic estimator can be interpreted as a first-order approximation to a maximum likelihood estimator for the lensing potential.

Refs. [23,30,46]. Adapting this notation, the expectation value of Eq. (7), averaged over realizations of CMB and lensing deflections (and assuming that both are Gaussian fields), becomes

$$\langle \hat{C}_L^{\kappa\kappa} \rangle = N_L^{(0)} + C_L^{\kappa\kappa} + N_L^{(1)} + \mathcal{O}[(C_L^{\kappa\kappa})^2]. \quad (8)$$

Nonlinear processes, such as nonlinear structure formation and multiple correlated deflections, introduce a small but detectable amount of non-Gaussianity to the lensing convergence [35,53–55]. In the limit of small density perturbations, the non-Gaussianity can be characterized by a lensing bispectrum. A nonzero lensing bispectrum changes the lensed temperature four-point function and introduces an additional bias term to Eq. (8),

$$\langle \hat{C}_L^{\kappa\kappa} \rangle = N_L^{(0)} + C_L^{\kappa\kappa} + N_L^{(1)} + N_L^{(3/2)} + \mathcal{O}[(C_L^{\kappa\kappa})^2]. \quad (9)$$

This new bias was first identified in Ref. [34]. We will now compare this theoretically derived term with measurements of a non-Gaussian bias in simulations.

III. SIMULATIONS

The general work flow for isolating a non-Gaussian bias is simple: we use a set of non-Gaussian convergence maps generated by ray tracing through an N-body simulation and another set of Gaussian convergence maps that has the same average power spectrum as the non-Gaussian maps. The same CMB realizations are lensed with both Gaussian and non-Gaussian convergence maps, which results in two sets of lensed CMB maps. We convolve these maps with a Gaussian beam before we add the same realizations of white measurement noise to both sets. We beam deconvolve the noisy maps before we apply the standard quadratic and four-point estimators. We then compare the results of the reconstructions between both sets and look for significant differences. In the following sections, we provide detailed descriptions and validations for each of these steps. The entire procedure is also illustrated in Fig. 1.

For all simulations (N-body and CMB), we use a standard ΛCDM cosmology with parameters, $H_0 = 72$ km/s/Mpc, $\Omega_m = 0.296$, $\sigma_8 = 0.786$, $w = -1$, $n_s = 0.96$ and $\Omega_b = 0.046$.

A. Convergence maps

We use a set of 10 240 non-Gaussian convergence maps that was obtained from ray tracing through snapshots of an N-body simulation. For a detailed description of their production we refer the reader to Ref. [55]. The underlying

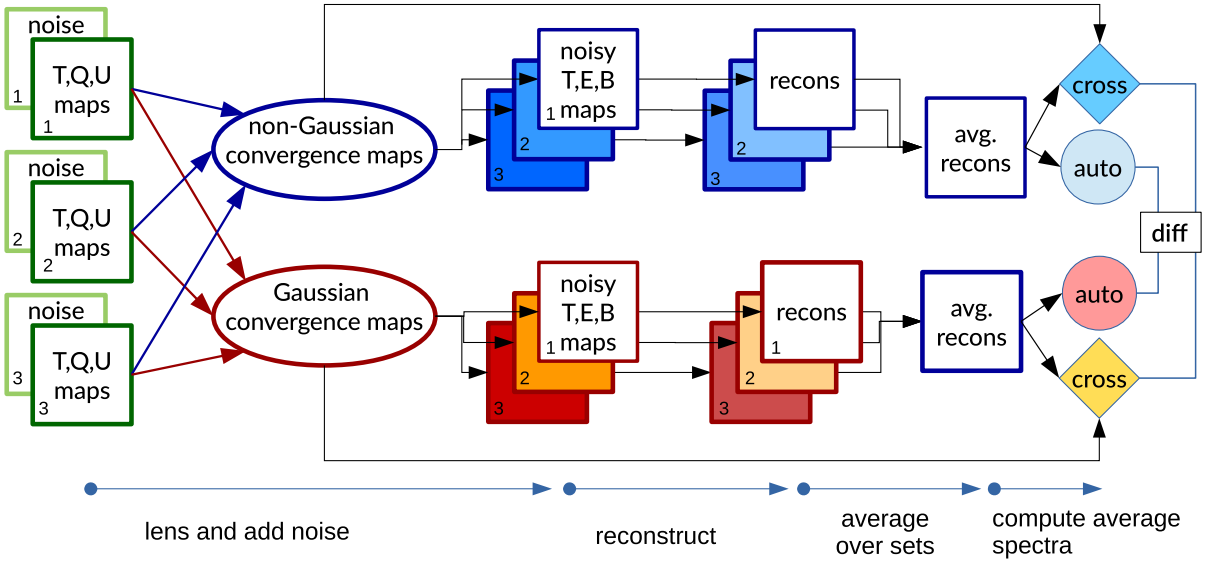

FIG. 1. Schematic outline of the simulation pipeline: squares and ellipses represent sets of 10 240 maps each. We start by generating 3 times 10 240 unlensed CMB and noise realizations. We then lens each CMB realization with both a Gaussian and a non-Gaussian convergence map. As we have 10 240 Gaussian and non-Gaussian convergence maps, each of them is used to lens three independent CMB realizations. This results in 2 times three sets of 10 240 lensed CMB maps. We convolve each of these maps with a Gaussian beam and add the same 3 times 10 240 noise realizations to the Gaussian and non-Gaussian maps. After beam deconvolution of the noisy maps, we run the standard quadratic estimator on every lensed CMB map. We average the reconstruction results over the maps that have been lensed with the same lensing convergence to beat down the reconstruction noise originating from the CMB sample variance. This leaves us with one set of 10 240 averaged Gaussian and another set of 10 240 averaged non-Gaussian reconstructed convergence maps. We compute the average power spectrum in each of these sets as well as the average cross-correlation with the true underlying convergence realizations. Any significant difference between the average power spectra of the non-Gaussian and Gaussian simulations is a non-Gaussian bias.

N-body simulation is based on the public Gadget-2 code [56], has a box size of 600 Mpc/h and is resolved by $N = 1024^3$ particles (corresponding to a mass resolution of $1.4 \times 10^{10}\ M_\odot/h$). The linear matter power spectrum for its initialization was computed with CAMB[3] [57] and initial conditions at $z = 100$ generated with N-GenIC. Snapshots were recorded between $z \approx 45$ and $z = 0$, a range which covers 99% of the growth corrected lensing kernel $W(\chi,\chi^*)D(z)$. Convergence maps were computed with LensTools [58] tracing 4096^2 light rays and calculating their deflections on 3 planes per box. This procedure does not assume that the deflection angle is small or that the light rays follow unperturbed geodesics. Different realizations of the convergence maps were produced by randomly rotating and shifting the potential planes [59]. The resulting maps are 12.25 deg^2 in size and resolved by 2048^2 pixels measuring 0.1025^2 arcmin^2. We refer to this simulation set as non-Gaussian or N-body lensing simulations. Their non-Gaussianity is a consequence not only of nonlinear structure formation in the N-body simulation but also of the multiple deflections along the lens planes. We do not measure or take into account the curl of the deflection field that is introduced by multiple deflection because we do not expect a significant bias from bispectra involving the curl component (see Appendix B for details).

We further produce a second set of 10 240 purely Gaussian convergence maps. These Gaussian simulations are generated by first measuring the average power spectrum of the non-Gaussian simulations and then drawing convergence realizations from a multivariate Gaussian with exactly this power spectrum.

In Fig. 2 we compare the average power spectra of the Gaussian and non-Gaussian simulation set to a theory power spectrum computed with the anisotropy solver CLASS[4] [60]. The missing power on the small-scale end, $L > 3000$, in the simulations is caused by the finite resolution of the N-body simulation [55]. On the large-scale end, the power is slightly suppressed because of the finite size of the simulation box. To allow for an accurate and unbiased detection of the non-Gaussian bias, we require an excellent agreement of the average power spectra in both simulation sets. Any significant difference between the power spectra could result in a false detection of a non-Gaussian bias. We find that the power spectrum of the Gaussian set agrees with the spectrum of the non-Gaussian simulations, as expected, within the sample variance (Fig. 3, red curve). We further compare the combined standard deviation of the average power in both simulation sets to the size of the bias predicted by BSS16 (Fig. 3 shaded region and blue dots). The comparison shows that the sample variance in the simulation sets is low enough to allow for a detection of a bias at the percent level (which corresponds to the magnitude predicted by BSS16).

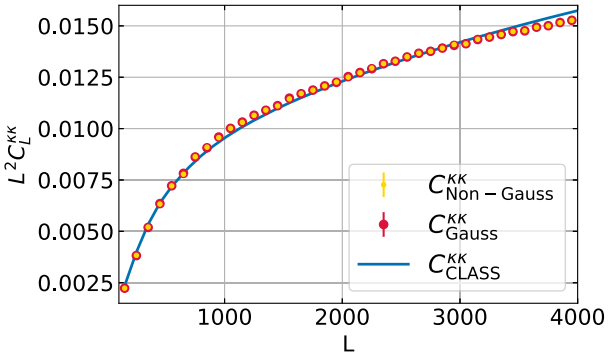

FIG. 2. Power spectra measured from 10 240 Gaussian (red) and ray-traced non-Gaussian (yellow) convergence maps closely follow the theory curve computed with CLASS (blue). Error bars corresponding to the standard deviation of the mean between the binned power spectra of the 10 240 maps are plotted but too small to be visible. For modeling nonlinear effects in the matter power spectrum CLASS uses a version of HALOFIT [61]. We use precision parameters tol_perturb_integration=1e-6, perturb_sampling_stepsize=0.01, k_min_tau0=0.002, k_max_tau0_over_l_max=10., halofit_k_per_decade=3000 and l_max_scalars=8000 to produce the theory curve. Missing power on small scales is owed to the finite resolution of the simulation.

To get a sense of the non-Gaussianity of the ray-traced convergence maps, we measure their skewness, $\langle\kappa(\mathbf{x})^3\rangle$ after smoothing them with a Gaussian kernel on different scales. The skewness is an integrated measure of the bispectrum.[5] By comparing the measurement with the theoretical prediction,

$$\langle\kappa(\mathbf{x})^3\rangle = \int_{\mathbf{l}}\int_{\mathbf{L}} W_R(L)W_R(l)W_R(|-\mathbf{L}-\mathbf{l}|) \times B^{\kappa\kappa\kappa}(L,l,|-\mathbf{L}-\mathbf{l}|) \quad (10)$$

$$W_R(l) = \exp\left(-l^2R^2/2\right), \quad (11)$$

we can determine the most suitable theoretical bispectrum model for computing the bias following BSS16.

We expect the bispectrum to have two contributions: one from nonlinear structure formation, where the convergence bispectrum is an integrated measure of the bispectrum of large-scales structure, and a second contribution from post-Born effects. In the squeezed limit these two contributions have opposite sign and partly cancel each other. We compare two different models for the convergence bispectrum induced by nonlinear structure formation; one in which we model the matter bispectrum in tree-level

[3] http://camb.info/
[4] http://class-code.net/

[5] Note that the quadratic estimator results in an additional skewness in the measured maps, i.e., the measured maps have nonzero skewness even if the underlying field is Gaussian [55]. We measure the skewness in the true noiseless convergence maps and not in the reconstructions as we are interested in quantifying the bispectrum introduced by nonlinear physics.

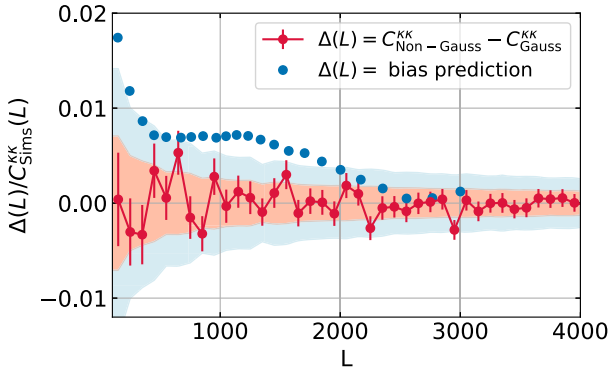

FIG. 3. For an accurate measurement of the nonlinear bias in the reconstructions it is crucial that the power spectra of the original, nonreconstructed, Gaussian and non-Gaussian simulations are consistent within their sample variance. We show that this is indeed the case by checking that their difference is consistent with zero (red dots, $\chi^2/\nu = 1.02$, with $\nu = 38$ degrees of freedom, corresponding to a p-value of 0.444). We further require the sample variance of both sets to be small enough to allow a detection of a bias at the percent level. To see this, we plot the combined sample variance of both sets as shaded regions ($\pm\sigma, \pm 2\sigma$) and compare it to the expected size of the nonlinear bias (blue dots).

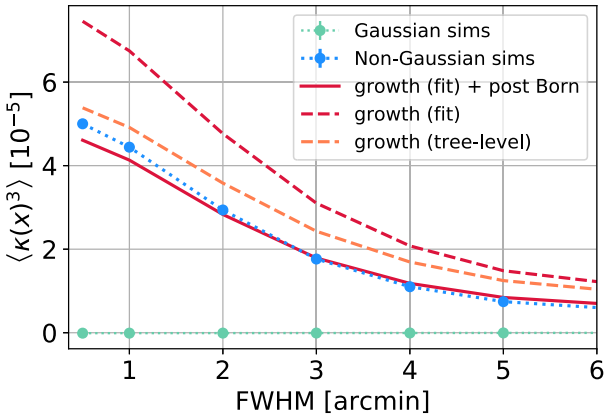

FIG. 4. Skewness measured on different scales provides some information on the bispectrum of the non-Gaussian convergence maps. We find that we can accurately model the skewness by assuming that the bispectrum consists of a nonlinear growth and post-Born induced contribution. The growth-induced part is best described by using a simulation-calibrated fit to the matter bispectrum [62]. The post-Born contribution is modeled following Ref. [35], and we find good agreement between theory and measurement only when this correction is included (solid red line). The convergence maps were smoothed with a Gaussian kernel with FWHMs indicated on the x-axis and filtered to exclude modes with $L > 4000$. For the theory curves we impose cutoffs at $k_{\rm min} = 0.0105[h/{\rm Mpc}]$ corresponding to the box size of the simulation and $k_{\rm max} = 50[h/{\rm Mpc}]$. Outside of these bounds we set the matter bispectrum (and matter power spectrum in the computation of the post-Born terms) to zero. From the comparison with a theory curve computed with $k_{\rm max} = 100[h/{\rm Mpc}]$, we find that the results are not sensitive to the $k_{\rm max}$ cutoff. The error bars correspond to the standard deviation of the mean and are smaller than the marker size.

perturbation theory[6] and one in which we use a simulation-calibrated fit to the matter bispectrum [62].[7] The convergence bispectrum from multiple deflections is modeled following Ref. [35]. The results of the skewness measurement together with the different theoretical models are shown in Fig. 4. We find that the theory curve computed from a combination of structure formation-induced and post-Born bispectra agrees well with the measurement on smoothing scales FWHM > 2 arcmin. On smaller scales we find a slight discrepancy, with the measurement lying above the theory prediction. We also find that simulation-calibrated fit to the matter bispectrum leads to better agreement with the simulation than the tree-level perturbation theory model. We use this best fitting model (red line in Fig. 4) in the following sections to compute the theoretical prediction for the nonlinear bias.

B. CMB simulations

We produce 3 times 10 240 unlensed CMB realization in temperature (T) and polarization (Q, U) based on power spectra computed with CAMB. This leaves us with three times more CMB realizations in each T, Q and U than convergence realizations. Consequently, we use each Gaussian and non-Gaussian convergence map to lens three independent CMB realizations. The reason for using the same lenses for a number of different background CMBs is that averaging over the lensing measurements from CMB maps lensed with the same convergence map reduces the reconstruction noise originating from the sample variance in the CMB. A lower reconstruction noise results in a more significant detection of the non-Gaussian bias.

The lensing algorithm used to produce the lensed CMB maps is described in detail in Ref. [63].[8] We apply a filter that removes modes with $L > 6000$ from the convergence maps prior to the lensing. This step is necessary for numerical stability and to remove unphysical effects that are caused by the finite resolution of the simulations. The power spectra of the lensed maps agree well with theory (as shown in Figs. 5 and 6): the relative difference between theory and simulations lies within $\pm 5\%$ for every bin in the

[6]In this model we replace the linear matter power spectrum by its HALOFIT counterpart.

[7]This bispectrum model has 9 free parameters which are assumed to be independent from cosmology and have been measured and fixed in Ref. [62]. It also depends on cosmological parameters through a direct appearance of σ_8 and indirectly by its dependence on the nonlinear scale and the nonlinear matter power spectrum. We adapt these quantities to agree with the cosmology of the simulation.

[8]We include terms up to fifth order in this algorithm.

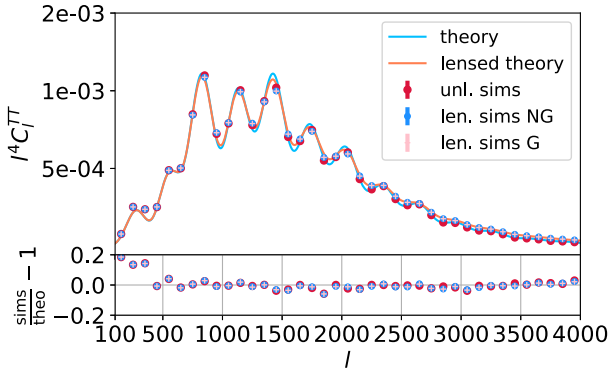

FIG. 5. Average power spectra of the lensed temperature maps in Gaussian and non-Gaussian simulation branches agree well with each other (lower panel). The agreement of lensed and unlensed realizations with the theory prediction is good except for large scales, where we find a significant deviation at $l < 500$. We exclude these scales from the reconstruction. The error bars corresponding to the standard deviation of the mean are too small to be resolved.

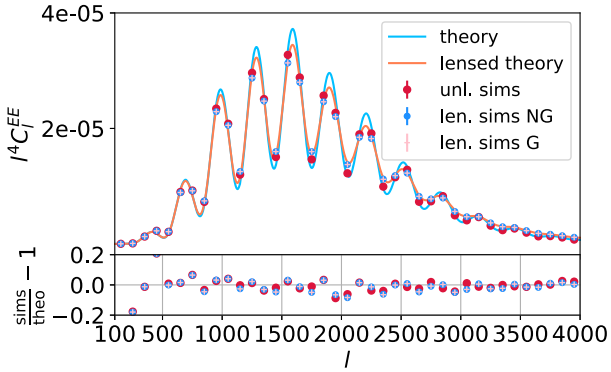

FIG. 6. Average power spectra of the unlensed and lensed polarization E-modes agree with the theory prediction except for large scales, where we find a significant deviation at $l < 500$. We exclude these scales from the reconstruction. Note that the error bars on the measured power spectra are smaller than the markers.

range $500 < L < 4000$. Crucial for the validity of this study, however, is the agreement between lensed Gaussian and lensed non-Gaussian CMB spectra. Comparing these two we find relative deviations of less than 0.3% per l-bin, which is consistent with the theoretical expectation [64].

After lensing, we convolve the lensed CMB maps with a Gaussian beam of width FWHM = 1 arcmin and add the same 3×10240 white noise realizations with a noise level of $\sigma_T = 1\ \mu\text{K} - \text{arcmin}$ in temperature and $\sigma_{\text{pol}} = \sqrt{2}\sigma_T$ in polarization to the Gaussian and non-Gaussian lensed CMB simulations. The noise configurations are chosen to roughly match prospective CMB surveys.

The entire procedure leaves us with 3×10240 mock CMB measurements for Gaussian and non-Gaussian convergence simulations each. Reconstructed convergence maps from Gaussian and non-Gaussian simulations differ only in their underlying convergence maps but have the same CMB and noise realizations.

C. Lensing reconstruction

We apply a quadratic estimator to all noisy, beam-deconvolved, lensed CMB maps in (TT) and (EB) to obtain noisy estimates of the underlying convergence fields. The reconstruction pipeline is described in detail in [14].

We filter scales with $l < 500$ from the CMB maps prior to reconstruction as Fig. 5 indicates some inconsistency between the power spectra of the lensed simulations and the theory power spectra on these scales. We also filter out any lensed CMB multipoles with $l > 4000$ which is a standard practice in CMB data analysis to exclude significant contamination from foregrounds and avoid foreground-related lensing biases. With $l = 4000$ we use a very optimistic scale; realistic CMB temperature data could already be contaminated by extragalactic foregrounds at lower multipoles [27]. Using $l < 3000$ reduces the theoretically predicted size of the bias by approx. a factor of 0.5 (see Fig. 14).

As three realizations in each branch have the same underlying convergence field, we average over their reconstructed power maps to recover a set of 10 240 power measurements corresponding to the 10 240 input maps in each simulation branch. These averaged measurements have reduced noise compared to measurements with only one CMB realization.

We proceed by computing the average power spectrum of the reconstructed lensing maps in the Gaussian and non-Gaussian branches for (TT) and (EB) reconstructions. By construction, we expect all lensing biases that are sensitive to the convergence power spectrum and the lensed or unlensed CMB power spectra [c.p. Eq. (8)] to be identical in the Gaussian and non-Gaussian simulations.[9] Because we are only interested in the difference of the reconstructed convergence power spectra, in which these biases cancel out, we do not compute and remove them. Apart from the auto power spectra, we also compute the average power in the cross-correlation between input maps and reconstructed maps. This cross-correlation is not an actual observable but can serve as a proxy for the non-Gaussian bias in cross-correlations with other tracers of large-scale structure. Also, measurements of the cross power are not affected by the $N^{(0)}$ and $N^{(1)}$ bias and have lower noise. The theory prediction for the bias in cross-correlations is $N_{\text{NG}}^{\text{cross}} \approx 1/2 N_{\text{NG}}^{\text{auto}}$ [34].

[9]The bispectrum of the convergence also changes the lensed CMB power spectra, but this is a subpercent effect and not detectable in our simulations [64].

IV. RESULTS

As discussed above we measure the non-Gaussian bias from the difference between

$$N^{\rm auto}_{\rm NG}(L) = \hat{C}^{\hat{\kappa}\hat{\kappa}}_{\rm NG}(L) - \hat{C}^{\hat{\kappa}\hat{\kappa}}_{\rm G}(L)$$
$$N^{\rm cross}_{\rm NG}(L) = \hat{C}^{\hat{\kappa}\kappa}_{\rm NG}(L) - \hat{C}^{\hat{\kappa}\kappa}_{\rm G}(L). \quad (12)$$

We start by examining this difference in the temperature based reconstructions.

In Fig. 7 we show the measured bias from cross-correlating reconstructions from (TT) with the input maps and plot the reconstructed power spectra for comparison. We plot the reconstruction results up to $L = 2000$. Including smaller scales than this is unlikely to increase the detection significance of the bias, given that the bias-to-signal and bias-to-noise ratios both decrease with increasing L. Figure 8 shows the measured bias in units of the signal. The theory prediction from BSS16 is plotted in light blue for comparison. A nonzero bias is detected with a significance of 5.21σ. The p-value of the data points assuming no bias is 0.0003. The p-value using the theoretical bias curve as null hypothesis amounts to 0.3932. A non-Gaussian bias is therefore detected with high statistical significance and is in agreement with the theoretical prediction. Measuring the covariance between the data points (Fig. 9, right panel) shows that the measurements in different bins can to good approximation be treated as uncorrelated.

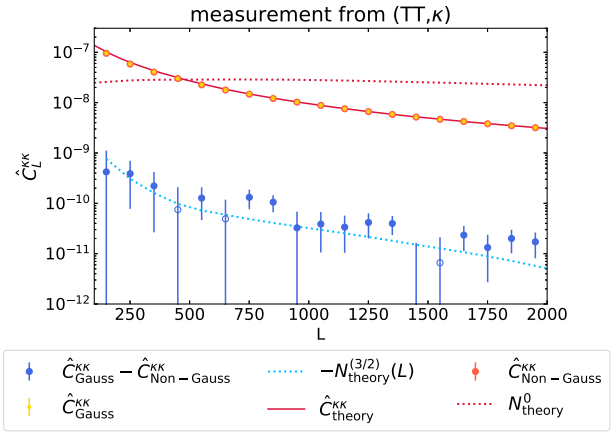

FIG. 7. Average power in the cross-correlation between convergence maps reconstructed from the lensed temperature maps and the input convergence field. The measured power follows the theory curve (red line) for both Gaussian and non-Gaussian simulations (yellow and orange points). We plot $C^{\kappa\kappa}_L$ computed with CLASS as theory prediction as neither $N^{(0)}$ nor $N^{(1)}$ lensing biases are expected to be present in the cross-correlation, and the $N^{(2)}$ bias should be greatly reduced by our usage of lensed CMB power spectra in the four-point estimator. The difference between the Gaussian and non-Gaussian reconstructions is shown as blue points (circles if they have negative sign). It is consistent with the theory prediction computed following BSS16 (light blue line).

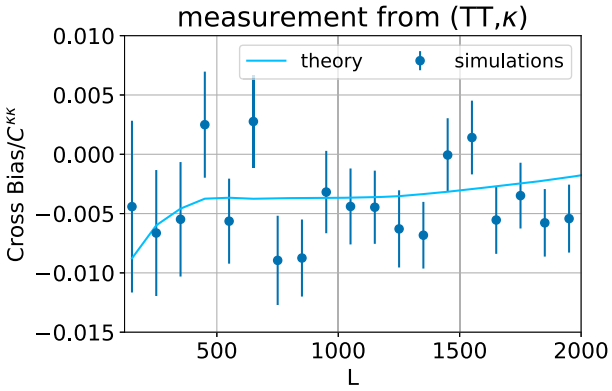

FIG. 8. We detect a non-Gaussian bias in the cross-correlation of temperature-based reconstructions and input maps with a significance of 5.21σ. For a null hypothesis of no bias we find a p-value of 0.0003. Using the theoretical bias curve computed with analytic expressions from BSS16 (plotted in light blue) as null hypothesis results in a p-value of 0.3932.

The bias in the auto power measured from (TT, TT) is detected with a lower significance of 2.84σ (Figs. 10 and 11). The p-value of the measurement for a null hypothesis of no bias is $p = 0.2660$. The null hypothesis can therefore not be rejected with high statistical significance. Using the theory prediction as null hypothesis, however, results in a even higher p-value of 0.7676. The errors in the auto spectrum are slightly correlated as the left panel of Fig. 9 shows. The correlation is below 20% for most of the bins and below 47% for all bins. This correlation is expected as the covariance of the four-point estimator is nondiagonal [23,65,66].

We do not find any indication for a non-Gaussian bias in the polarization-based reconstruction (Fig. 12). This agrees with the intuition gained from its functional form: additional angular dependencies (as compared to the bias in temperature-based reconstruction) reduce the support of the contributing integrals (see Appendix A and Ref. [34]).

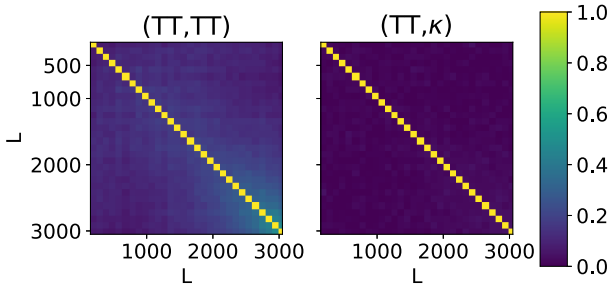

FIG. 9. Covariance matrices of the power spectrum measurements in units of the variance. The error bars in the measurement of the power in the cross-correlation are to good approximation uncorrelated. The measurements of the auto power show some expected degree of correlation between the bins.

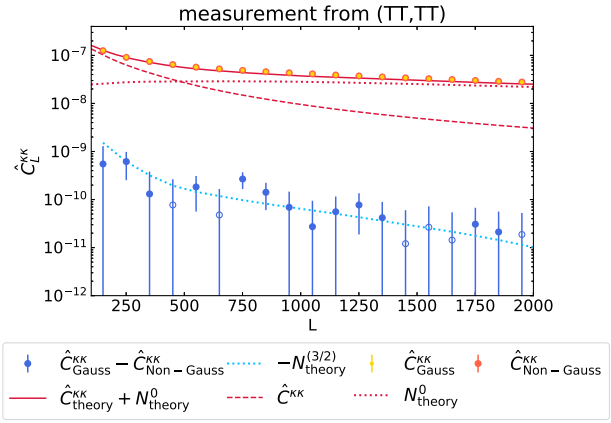

FIG. 10. Measured non-Gaussian bias in the CMB lensing power spectrum measurement from the temperature four-point function. The reconstruction agrees well with the theory prediction (red curve) which we model as a sum of convergence power $C^{\kappa\kappa}_L$ and $N^{(0)}$ lensing bias. The measured non-Gaussian bias is consistent with the theory prediction of BSS16, but the null hypothesis of no bias cannot be excluded with high statistical significance.

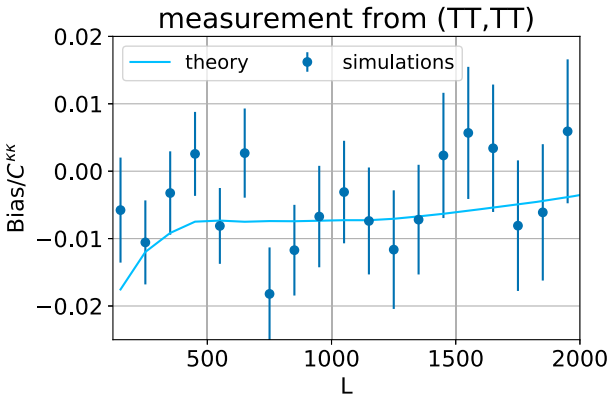

FIG. 11. Non-Gaussian bias in temperature-based CMB lensing power spectrum measurements in units of the signal. The points are consistent with the theory predictions, and the bias is detected with a significance of 2.84σ. We find a p-value of 0.2660 for a no-bias null hypothesis and a p-value of 0.7676 when using the theory prediction as null hypothesis.

Again, we find correlations between the data points in the auto power measurement (Fig. 13).

V. DISCUSSION

By comparing lensing measurements from CMB simulations lensed with Gaussian and non-Gaussian convergence fields, we find strong indication for the existence of a non-Gaussian bias to CMB lensing measurements from temperature data. The bias is at the 1% level, which agrees with the theoretical prediction for a bispectrum-induced bias of Ref. [34] if we take into account two sources for the

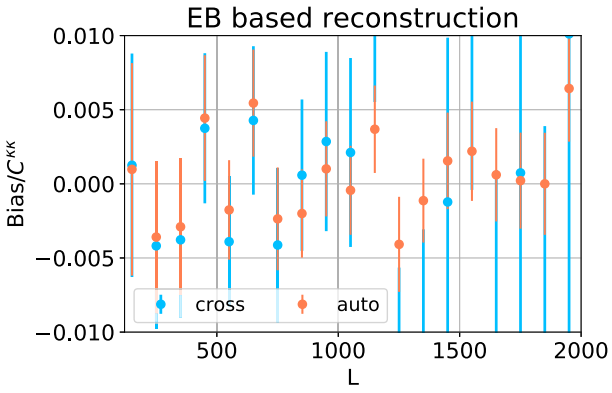

FIG. 12. Non-Gaussian bias to CMB lensing measurements from (EB, EB) is consistent with zero in both auto and cross-correlations. This seems in agreement with BSS16, who have not numerically evaluated the bias expressions for polarization, but argue that they are expected to be smaller than in temperature due to their additional angular dependencies.

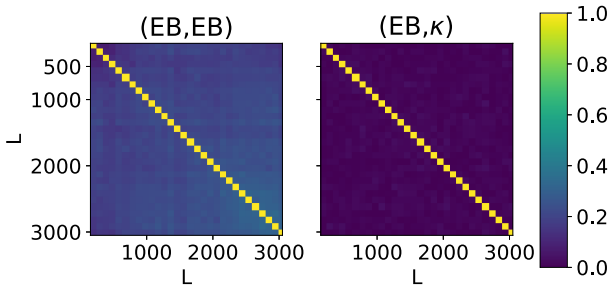

FIG. 13. Covariance matrices of the power spectrum measurements from (EB, EB) in units of the variance. The error bars in the measurement of the power in the cross-correlation are to good approximation uncorrelated. The measurements of the auto power show some expected degree of correlation between the bins.

lensing bispectrum, nonlinear structure formation and multiple correlated deflections.

By measuring the bias in the cross-correlations of reconstructed lensing maps with the true underlying lensing fields, we detect the theoretically predicted bias in the simulations at the 5σ significance level. We detect the non-Gaussian bias in the autocorrelation with a significance of $\sim 3\sigma$. The measured bias in power spectrum measurements from a combination of E- and B-mode polarization, (EB, EB), is consistent with zero. We note that lensing B-modes at intermediate scales are more sensitive to smaller scales in the deflection field than lensed E-modes or temperature. A nonzero bias in EB, EB could therefore be present in real data if it was generated by scales that are not accurately modeled in the simulation due to its finite resolution.

We point out that our results have been independently confirmed by Ref. [38], who use a completely different simulation set on the full sky.

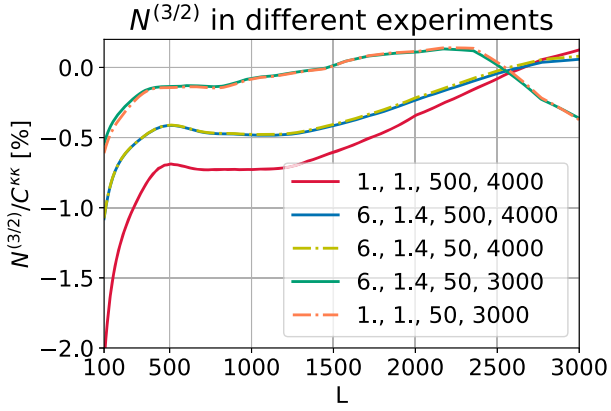

FIG. 14. Bispectrum-induced $N^{(3/2)}$ bias in temperature-based measurements for different experimental configurations. Different from the results shown in BSS16, these curves are derived from a lensing bispectrum model that includes post-Born correction and uses a simulation-calibrated fitting formula to the matter bispectrum. The lines are labeled by noise in $\mu K-$ arcmin, beam FWHM in arcmin, $l_{\rm min}$ and $l_{\rm max}$. The size of the bias is sensitive to the maximal, signal-dominated CMB scale that is used in the reconstruction.

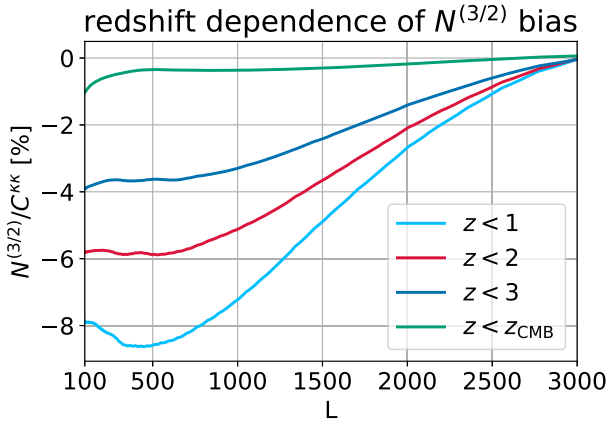

FIG. 15. Relative size of the non-Gaussian bias increases if we only consider lenses at low redshifts $z_{\rm max}<z_{\rm CMB}$. This is a consequence of the different redshift scalings of the competing terms in the lensing bispectrum from nonlinear structure formation and post-Born effects. To illustrate this we ignore any contributions to the lensing bispectrum and power spectrum with $z>z_{\rm max}$ and plot the ratio of the resulting cross bias to the power spectrum. These results suggest that the non-Gaussian bias could be more important for measurements of cross-correlations of CMB lensing with low-redshift tracers. We note that the curves shown here are still preliminary and should be seen as a motivation to investigate the bias on cross-correlations in future work (Böhm *et al.* in prep). Note that the curves show the bias from gradient modes only and do not take into account lensing by curl modes.

The good agreement between the simulations and theory suggests that the assumptions that entered into the theory calculation are valid and that we can rely on it to make predictions for different experiments. Theoretical bias predictions for different experimental configuration are shown in Fig. 14.

The non-Gaussian bias is likely to affect lensing measurements from CMB experiments that are dominated by temperature reconstruction. This includes current and upcoming experiments such as AdvACT [16] and Simons Observatory [67]. An uncorrected non-Gaussian bias at the percent level degrades the accuracy with which these experiments can measure cosmological parameters. The non-Gaussian bias is unlikely to affect experiments that are polarization dominated, such as the ground-based SPT-3G [18] and CMB-S4 experiments [19] and space-based missions like LiteBird [68] or Pico [69].

It is further important to note that the smallness of the bias is a consequence of a somewhat coincidental cancellation: The bias is mostly sensitive to elongated bispectrum configurations. For these shapes, the bispectra from nonlinear structure formation and multiple correlated deflections have opposite sign. The fact that they are in addition of similar magnitude is only true for sources at high redshifts. If we consider sources at low redshifts or restrict contributions to the bispectra to low redshifts, we expect this cancellation to be much less efficient. The bias could therefore be more important in cross-correlations of CMB lensing with low-redshift tracers (Böhm *et al.* in prep., Ref. [70]). We illustrate this by plotting preliminary results for the non-Gaussian bias (corrected by a factor of $1/2$ as it is expected for cross-correlations) in units of the CMB lensing signal when only allowing lenses at low redshift to contribute to both ($z_{\rm max}\leq z_{\rm CMB}, z_{\rm source}=z_{\rm CMB}$) in Fig. 15. These results suggest that the bias could be on the order of several percent for cross-correlation measurements from temperature data. With these measurements getting most of their signal from high multipoles where the TT estimator performs best, this could make this bias relevant for most future wide-field surveys. We caution at this point that Fig. 15 should only be seen as a motivation to investigate the nonlinear bias for cross-correlations. By setting all contribution to the post-Born bispectrum above a certain redshift to zero, it becomes negligible. For realistic cross-correlations the expression for the post-Born bispectrum is more complicated, and its contribution to the bias could be more important.

The results shown in this work only apply to power spectrum estimates with a quadratic estimator, but similar biases could arise for alternative estimators if they are derived under the assumption of a Gaussian deflection field.

Recently, Ref. [52] pointed out that a shear estimator [50,51] is to good approximation robust against contamination from isotropic foregrounds (at the cost of lower signal to noise in the reconstruction). The fact that we find no bias in the reconstruction from EB, EB, for which the quadratic estimator corresponds to a shear estimator, suggests that a shear-only estimator could also be less

sensitive to the nonlinear bias [this can also be seen analytically because the shear estimator has an additional angular dependence, which should lead to additional cancellations in the bias integrals in Eq. (A1)]. This possibility could be easily tested on simulations and could be explored in future work. Another mitigation strategy that could be explored is removing small scales (e.g., with $l \gtrsim 2000$) in the observed lensed CMB that enters the gradient leg of the quadratic estimator. This strategy is motivated by the fact that the bias scales strongly with the maximum CMB scale used in the reconstruction. It has been shown to reduce similar biases in cluster lensing [71] as well as biases originating from foreground contamination [33], while only losing little signal-to-noise. Even though all these paths are interesting and should be used as independent checks, we stress that the results obtained in this work suggest that the bias can be theoretically modeled and subtracted. The cosmology dependence of the bias, specifically on σ_8, might require iterative estimation and raises the question if the bias could be used as a signal on its own. Given its smallness compared to the lensing reconstruction noise, this does not seem very promising. The presence of this bias, however, should be taken as a motivation for the design of estimators that implicitly take the non-Gaussian structure of the lensing convergence into account, thus exploring its full information content.

ACKNOWLEDGMENTS

We thank Dominic Beck and Giulio Fabbian for cross-checking some of the results with their simulations and Antony Lewis and Geraint Pratten for sharing their numerical implementations of post-Born terms. V. B. thanks Emmanuel Schaan, Dominic Beck and Giulio Fabbian for useful discussions. B. D. S. was supported by an STFC Ernest Rutherford Fellowship and an Isaac Newton Trust Early Career Grant. J. L. is supported by an NSF Astronomy and Astrophysics Postdoctoral Fellowship under Grant No. AST-1602663. J. C. H. is supported by the Friends of the Institute for Advanced Study. M. S. was supported by the Bezos fund. This work used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by NSF Grant No. ACI-1053575. This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. Part of this work used computational resources at the Max Planck Computing and Data Facility (MPCDF).

APPENDIX A: ANALYTIC PREDICTION FOR A BISPECTRUM-INDUCED CMB LENSING BIAS

All CMB lensing analyses to date assume that the lensing convergence is a Gaussian field. However, nonlinear structure formation and multiple correlated lenses introduce a small, but detectable amount of non-Gaussianity [35,53–55]. In the limit of small density perturbations, the non-Gaussian structure can be characterized by a hierarchy of connected correlation functions. To lowest order, the lensing convergence acquires a bispectrum.

The lensing bispectrum introduces an additional term to the standard four-point estimator [compare Eq. (9)]. This new bias was first identified in Ref. [34] (BSS16). Its name follows from the naming convention for CMB lensing biases, where biases are labeled by their power in the lensing power spectrum. The $N^{3/2}$ bias arises because the lensing bispectrum changes the lensed temperature four-point function.

BSS16 estimated that the $N^{3/2}$ bias could change the measured lensing power spectrum in temperature-based CMB lensing analyses at the percent level.

The estimation of the size of the bias in BSS16 is based on the numerical evaluation of analytically derived expressions. This evaluation relies on a number of assumptions:

(1) The lensing bispectrum contributes to the lensed temperature four-point function with 8 terms. Due to the complicated structure of these terms (they involve six-dimensional coupled integrals over reconstruction weights g, the lensing bispectrum and CMB power spectra), only two of these terms were evaluated. These two terms were chosen because they factor maximally under the reconstruction weights (one of them can even be split into a product of 3 two-dimensional integrals). Their structure suggests that these terms are the dominant contributions to the bispectrum-induced bias.

For temperature-only reconstruction, the two terms read

$$\begin{aligned} N_1^{(3/2)}(L) &= -4A_L^2 S_L \int_{\mathbf{l}_1,\mathbf{l}} g_{\mathbf{l}_1,\mathbf{L}}[\mathbf{l}\cdot(\mathbf{l}_1-\mathbf{l})] \\ &\quad \times [\mathbf{l}\cdot(\mathbf{L}-(\mathbf{l}_1-\mathbf{l}))] \\ &\quad \times C_l^{TT} B_\phi[\mathbf{l}_1-\mathbf{l}, \mathbf{L}-(\mathbf{l}_1-\mathbf{l}), -\mathbf{L}] \\ N_2^{(3/2)}(L) &= 4A_L^2 S_L \int_{\mathbf{l}_1,\mathbf{l}} g_{\mathbf{l}_1,\mathbf{L}}(\mathbf{l}_1\cdot\mathbf{l})[\mathbf{l}_1\cdot(\mathbf{L}-\mathbf{l})] \\ &\quad \times C_{l_1}^{TT} B_\phi(\mathbf{l}, \mathbf{L}-\mathbf{l}, -\mathbf{L}), \end{aligned} \tag{A1}$$

with

$$S_L = \int_{\mathbf{l}_2} g_{\mathbf{l}_2,\mathbf{L}}(\mathbf{l}_2\cdot\mathbf{L}) C_{l_2}^{TT} \approx \frac{1}{2} A_L^{-1}. \tag{A2}$$

For polarization-based reconstruction, the structure of the terms is similar, but with additional angular dependencies (see Ref. [34] for details).

In cross-correlations, all other terms vanish and the bias depends only on the two terms above.

(2) The evaluation in BSS16 only considered nonlinear structure formation as a source of the lensing bispectrum. Recently, Ref. [35] pointed out that an additional lensing bispectrum arises from multiple correlated lensing deflections. The effect of multiple deflections is commonly ignored in the Born approximation. Both effects, post-Born corrections and nonlinear structure formation, lead to lensing bispectra of the same order of magnitude but for certain triangle configurations of opposite sign.

(3) The modeling of the bispectrum from nonlinear structure formation in BSS16 relied on a tree-level perturbation theory, which breaks down on small scales (and the bias was shown to be sensitive to replacing the linear matter power spectrum by its nonlinear (HALOFIT [61]) counterpart in the matter bispectrum model).

(4) The theoretical modeling of the bias relied on a Taylor series expansion of the lensed CMB in the deflection angle and thus on the assumption of small deflection angles.

(5) Whereas BSS16 provide theory expressions for the bias to lensing measurements from temperature and polarization, they have only numerically evaluated the bias for temperature-based reconstruction. The expressions for polarization have additional angular dependencies which introduce oscillations that makes their evaluation numerically very challenging. These oscillations should also average out in the integrations and result in smaller bias terms in polarization than in temperature.

In this work, we use an updated analytical prediction for the bias. We use the expressions given in BSS16 but use a different bispectrum model that includes the bispectrum from post-Born effects and uses an extended, simulation-calibrated, semianalytic model for the matter bispectrum [62]. Results with this new bispectrum model are shown in Fig. 14 for different experimental setups and plotted together with the measurement of the bias in Sec. IV. For a CMB-S4-like experiment (red curve in Fig. 14), we find a cumulative bias over noise of $\sim 3\sigma$ in lensing reconstruction from temperature only (assuming $f_{\rm sky} = 0.5$ and a bin width of $\Delta L = 100$). The maximum bias-over-noise per L-bin in this configuration is 1 in the lowest bin (centered on $L = 150$).

APPENDIX B: NONLINEAR BIAS FROM BISPECTRA INVOLVING THE CURL OF THE LENSING DEFLECTION

Allowing for multiple deflections introduces an additional degree of freedom, ω, to the linear mapping between the lensed and unlensed image of a source, which describes a rotation of the image [72]. With this additional degrees of freedom, the lensing deflection angle is no longer a pure gradient field but acquires an additional curl component,

$$\alpha(\mathbf{x}) = \nabla\phi(\mathbf{x}) + *\nabla\Omega(\mathbf{x}), \tag{B1}$$

sourced by the curl potential Ω [73]. We use an $*$ to denote a rotation by 90 degrees and, for notational simplicity, also abbreviate the combination of rotation and scalar product, $\cdot*$, in the following by $*$.

Being second order in the gravitational potential, the rotation is suppressed compared to the first-order convergence and shear distortions to the image. We thus expect the largest bias that involves the curl potential to be sourced by a "cross" bispectrum of the form $B^{\Omega,\phi,\phi}(\mathbf{L},\mathbf{l},-\mathbf{L}-\mathbf{l})$ [35]. The curl potential can be treated in complete analogy to the scalar lensing potential ϕ. For example, when expressing the effect of lensing on the CMB in terms of a small perturbation to the unlensed CMB, we can write [74]

$$\tilde{T} = T + \delta_\Omega T + \delta_\phi T + \delta_\Omega^2 T + \delta_\phi^2 T + \mathcal{O}(\phi^3, \Omega^3). \tag{B2}$$

Adapting the flat sky approximation, the first two terms are given in harmonic space by

$$\delta_\Omega T(\mathbf{l}) = \int_{\mathbf{l}'} [\mathbf{l}' * (\mathbf{l}' - \mathbf{l})] T(\mathbf{l}') \Omega(\mathbf{l} - \mathbf{l}') \tag{B3}$$

$$\delta_\phi T(\mathbf{l}) = \int_{\mathbf{l}'} [\mathbf{l}' \cdot (\mathbf{l}' - \mathbf{l})] T(\mathbf{l}') \phi(\mathbf{l} - \mathbf{l}'). \tag{B4}$$

Using this perturbative framework to model the lensed temperature four-point function, Ref. [34] shows that the two dominant terms in the $N^{(3/2)}$ bias are sourced by contractions of the following expectation values over ϕ and T [34],[10]

$$\begin{aligned} N_1^{(3/2)}[\phi^3] &\leftarrow \langle \delta_\phi T \delta_\phi T \delta_\phi T' T' \rangle \\ N_2^{(3/2)}[\phi^3] &\leftarrow \langle \delta_\phi T T \delta_\phi^2 T' T' \rangle. \end{aligned} \tag{B5}$$

A bias sourced by the cross bispectrum $B^{\Omega,\kappa,\kappa}(\mathbf{L},\mathbf{l},-\mathbf{L}-\mathbf{l})$ (we refer to it as $\tilde{N}^{(3/2)}$) should therefore be dominated by contractions of the following expectation values:

[10]By contractions we mean the terms that arise from taking the expectation value over unlensed CMB realizations. Assuming that the unlensed CMB is Gaussian, each expectation value can be split into a sum of three terms. We only consider the bias arising from one of these three terms, which BSS16 identified as the dominant one. Also, for readability, we do not write the symmetry factors that arise from permutations of $T, \delta T$ and $\delta^2 T$ that leave the result invariant.

$$\tilde{N}_1^{(3/2)}[\phi^2\Omega] \leftarrow \langle \delta_\phi T \delta_\phi T \delta_\Omega T' T' \rangle_{1a} + \langle \delta_\Omega T \delta_\phi T \delta_\phi T' T' \rangle_{1b}$$
$$\tilde{N}_2^{(3/2)}[\phi^2\Omega] \leftarrow \langle \delta_\Omega T T \delta_\phi^2 T' T' \rangle_2. \quad \text{(B6)}$$

The expressions for the dominant contractions arising from $1a$ and 2 are identical to the auto bias [Eq. (A1)], but with B^{ϕ^3} replaced by $B^{\phi^2\Omega}$ and S_L replaced by,

$$S_L^\times = \int_{\mathbf{l}_2} g_{\mathbf{l}_2,\mathbf{L}}(\mathbf{l}_2 * \mathbf{L}) C_{l_2}^{TT} = 0. \quad \text{(B7)}$$

This integral vanishes because the integrand is uneven under the angular integration. The remaining dominant contraction from $1b$ is of the form

$$\tilde{N}_{1b}^{(3/2)}(L) = -4A_L^2 S_L \int_{\mathbf{l}_1,\mathbf{l}} g_{\mathbf{l}_1,\mathbf{L}}[\mathbf{l} * \mathbf{l}_1] \times [\mathbf{l} \cdot (\mathbf{L} - (\mathbf{l}_1 - \mathbf{l}))] \times C_l^{TT} B_{\phi,\phi,\Omega}[\mathbf{l}_1 - \mathbf{l}, \mathbf{L} - (\mathbf{l}_1 - \mathbf{l}), -\mathbf{L}]. \quad \text{(B8)}$$

Because of the mixing of sines and cosines in the angular integrations in Eq. (B8), we expect this contribution to be strongly suppressed compared to the corresponding term in the bias from the auto bispectrum.

This short calculation suggests that biases from bispectra involving the curl component are likely to be negligible for current and upcoming CMB experiments.

[1] A. Lewis and A. Challinor, Phys. Rep. **429**, 1 (2006).
[2] D. Hanson, A. Challinor, and A. Lewis, Gen. Relativ. Gravit. **42**, 2197 (2010).
[3] J. Lesgourgues, L. Perotto, S. Pastor, and M. Piat, Phys. Rev. D **73**, 045021 (2006).
[4] B. D. Sherwin *et al.*, Phys. Rev. Lett. **107**, 021302 (2011).
[5] K. M. Smith, O. Zahn, and O. Doré, Phys. Rev. D **76**, 043510 (2007).
[6] C. M. Hirata, S. Ho, N. Padmanabhan, U. Seljak, and N. A. Bahcall, Phys. Rev. D **78**, 043520 (2008).
[7] S. Das *et al.*, Phys. Rev. Lett. **107**, 021301 (2011).
[8] A. van Engelen *et al.*, Astrophys. J. **756**, 142 (2012).
[9] D. Hanson *et al.*, Phys. Rev. Lett. **111**, 141301 (2013).
[10] P. A. R. Ade *et al.* (Polarbear Collaboration), Phys. Rev. Lett. **113**, 021301 (2014).
[11] P. A. R. Ade *et al.* (BICEP2 and Keck Array Collaborations), Astrophys. J. **833**, 228 (2016).
[12] K. T. Story *et al.*, Astrophys. J. **810**, 50 (2015).
[13] P. A. R. Ade *et al.* (Planck Collaboration), Astron. Astrophys. **594**, A15 (2016).
[14] B. D. Sherwin *et al.*, Phys. Rev. D **95**, 123529 (2017).
[15] G. Simard *et al.*, Astrophys. J. **860**, 137 (2018).
[16] S. W. Henderson *et al.*, J. Low Temp. Phys. **184**, 772 (2016).
[17] A. Suzuki *et al.*, J. Low Temp. Phys. **184**, 805 (2016).
[18] B. A. Benson *et al.*, Proc. SPIE Int. Soc. Opt. Eng. **9153**, 91531P (2014).
[19] K. N. Abazajian *et al.*, arXiv:1610.02743.
[20] W. Hu, Astrophys. J. **557**, L79 (2001).
[21] W. Hu and T. Okamoto, Astrophys. J. **574**, 566 (2002).
[22] M. Kesden, A. Cooray, and M. Kamionkowski, Phys. Rev. D **67**, 123507 (2003).
[23] D. Hanson, A. Challinor, G. Efstathiou, and P. Bielewicz, Phys. Rev. D **83**, 043005 (2011).
[24] D. Hanson, G. Rocha, and K. Górski, Mon. Not. R. Astron. Soc. **400**, 2169 (2009).
[25] D. Hanson, A. Lewis, and A. Challinor, Phys. Rev. D **81**, 103003 (2010).
[26] Y. Fantaye, C. Baccigalupi, S. M. Leach, and A. P. S. Yadav, J. Cosmol. Astropart. Phys. 12 (2012) 017.
[27] A. van Engelen, S. Bhattacharya, N. Sehgal, G. P. Holder, O. Zahn, and D. Nagai, Astrophys. J. **786**, 13 (2014).
[28] S. J. Osborne, D. Hanson, and O. Doré, J. Cosmol. Astropart. Phys. 3 (2014) 024.
[29] S. Ferraro and J. C. Hill, Phys. Rev. D **97**, 023512 (2018).
[30] A. Lewis, A. Challinor, and D. Hanson, J. Cosmol. Astropart. Phys. 03 (2011) 018.
[31] T. Namikawa, D. Hanson, and R. Takahashi, Mon. Not. R. Astron. Soc. **431**, 609 (2013).
[32] T. Namikawa and R. Takahashi, Mon. Not. R. Astron. Soc. **438**, 1507 (2014).
[33] M. S. Madhavacheril and J. C. Hill, Phys. Rev. D **98**, 023534 (2018).
[34] V. Böhm, M. Schmittfull, and B. D. Sherwin, Phys. Rev. D **94**, 043519 (2016).
[35] G. Pratten and A. Lewis, J. Cosmol. Astropart. Phys. 08 (2016) 047.
[36] A. Petri, Z. Haiman, and M. May, Phys. Rev. D **95**, 123503 (2017).
[37] G. Fabbian, M. Calabrese, and C. Carbone, J. Cosmol. Astropart. Phys. 02 (2018) 050.
[38] D. Beck, G. Fabbian, and J. Errard, Phys. Rev. D **98**, 043512 (2018).
[39] B. Jain, U. Seljak, and S. White, Astrophys. J. **530**, 547 (2000).
[40] D. Munshi, P. Valageas, L. van Waerbeke, and A. Heavens, Phys. Rep. **462**, 67 (2008).
[41] M. Loverde and N. Afshordi, Phys. Rev. D **78**, 123506 (2008).
[42] V. Assassi, M. Simonović, and M. Zaldarriaga, J. Cosmol. Astropart. Phys. 11 (2017) 054.
[43] P. Lemos, A. Challinor, and G. Efstathiou, J. Cosmol. Astropart. Phys. 05 (2017) 014.
[44] M. Kilbinger, C. Heymans, M. Asgari, S. Joudaki, P. Schneider, P. Simon, L. Van Waerbeke, J. Harnois-Déraps,

H. Hildebrandt, F. Köhlinger, K. Kuijken, and M. Viola, Mon. Not. R. Astron. Soc. **472**, 2126 (2017).
[45] M. Zaldarriaga and U. Seljak, Phys. Rev. D **59**, 123507 (1999).
[46] E. Anderes, Phys. Rev. D **88**, 083517 (2013).
[47] C. M. Hirata and U. Seljak, Phys. Rev. D **67** (2003).
[48] J. Carron and A. Lewis, Phys. Rev. D **96** (2017).
[49] M. Millea, E. Anderes, and B. D. Wandelt, arXiv:1708.06753.
[50] M. Bucher, C. S. Carvalho, K. Moodley, and M. Remazeilles, Phys. Rev. D **85**, 043016 (2012).
[51] H. Prince, K. Moodley, J. Ridl, and M. Bucher, J. Cosmol. Astropart. Phys. 01 (2018) 034.
[52] E. Schaan and S. Ferraro, arXiv:1804.06403.
[53] T. Namikawa, Phys. Rev. D **93**, 121301 (2016).
[54] G. Marozzi, G. Fanizza, E. Di Dio, and R. Durrer, J. Cosmol. Astropart. Phys. 09 (**2016**) 028.
[55] J. Liu, J. C. Hill, B. D. Sherwin, A. Petri, V. Böhm, and Z. Haiman, Phys. Rev. D **94**, 103501 (2016).
[56] V. Springel, Mon. Not. R. Astron. Soc. **364**, 1105 (2005).
[57] A. Lewis, A. Challinor, and A. Lasenby, Astrophys. J. **538**, 473 (2000).
[58] A. Petri, Astron. Comput. **17**, 73 (2016).
[59] A. Petri, Z. Haiman, and M. May, Phys. Rev. D **93**, 063524 (2016).
[60] D. Blas, J. Lesgourgues, and T. Tram, J. Cosmol. Astropart. Phys. 07 (2011) 034.
[61] R. Takahashi, M. Sato, T. Nishimichi, A. Taruya, and M. Oguri, Astrophys. J. **761**, 152 (2012).
[62] H. Gil-Maret alín, C. Wagner, F. Fragkoudi, R. Jimenez, and L. Verde, J. Cosmol. Astropart. Phys. 02 (2012) 047.
[63] T. Louis, S. Næss, S. Das, J. Dunkley, and B. Sherwin, Mon. Not. R. Astron. Soc. **435**, 2040 (2013).
[64] A. Lewis and G. Pratten, J. Cosmol. Astropart. Phys. 12 (2016) 003.
[65] M. M. Schmittfull, A. Challinor, D. Hanson, and A. Lewis, Phys. Rev. D **88**, 063012 (2013).
[66] J. Peloton, M. Schmittfull, A. Lewis, J. Carron, and O. Zahn, Phys. Rev. D **95**, 043508 (2017).
[67] P. Ade *et al.* (Simons Observatory Collaboration), arXiv:1808.07445.
[68] A. Suzuki *et al.*, arXiv:1801.06987.
[69] B. M. Sutin *et al.*, *Space Telescopes and Instrumentation 2018: Optical, Infrared, and Millimeter Wave* Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 10698 (2018), 106984F.
[70] P. M. Merkel and B. M. Schäfer, Mon. Not. R. Astron. Soc. **471**, 2431 (2017).
[71] W. Hu, S. DeDeo, and C. Vale, New J. Phys. **9**, 441 (2007).
[72] A. Stebbins, arXiv:astro-ph/9609149.
[73] C. M. Hirata and U. Seljak, Phys. Rev. D **68**, 083002 (2003).
[74] A. Cooray, M. Kamionkowski, and R. R. Caldwell, Phys. Rev. D **71**, 123527 (2005).