

This is an Open Access document downloaded from ORCA, Cardiff University's institutional repository:<https://orca.cardiff.ac.uk/id/eprint/174239/>

This is the author's version of a work that was submitted to / accepted for publication.

Citation for final published version:

Darwish, Omar, Belkner, Sebastian, Legrand, Louis, Carron, Julien and Fabbian, Giulio 2024. Non-Gaussian deflections in iterative optimal CMB lensing reconstruction. *Physical Review D (particles, fields, gravitation, and cosmology)* 110 (10) , 103520. 10.1103/physrevd.110.103520

Publishers page: <https://doi.org/10.1103/physrevd.110.103520>

Please note:

Changes made as a result of publishing processes such as copy-editing, formatting and page numbers may not be reflected in this version. For the definitive version of this publication, please refer to the published source. You are advised to consult the publisher's version if you wish to cite this paper.

This version is being made available in accordance with publisher policies. See <http://orca.cf.ac.uk/policies.html> for usage policies. Copyright and moral rights for publications made available in ORCA are retained by the copyright holders.



Non-Gaussian deflections in iterative optimal CMB lensing reconstruction

Omar Darwish,^{1,*} Sebastian Belkner,^{1,†} Louis Legrand,^{1,2} Julien Carron,¹ and Giulio Fabbian^{3,4,5,6}

¹*Université de Genève, Département de Physique Théorique et CAP, 24 Quai Ansermet, CH-1211 Genève 4, Switzerland*

²*Instituto de Física Teórica da Universidade Estadual Paulista and ICTP South American Institute for Fundamental Research, R. Dr. Bento Teobaldo Ferraz, 271, Bloco II, Barra-Funda - São Paulo/SP, Brasil*

³*Institute of Astronomy, Madingley Road, Cambridge CB3 0HA, UK*

⁴*Kavli Institute for Cosmology Cambridge, Madingley Road, Cambridge CB3 0HA, UK*

⁵*School of Physics and Astronomy, Cardiff University, The Parade, Cardiff, Wales CF24 3AA, United Kingdom*

⁶*Center for Computational Astrophysics, Flatiron Institute, 162 Fifth Avenue, New York, NY, 10010, USA*

The gravitational lensing signal from the Cosmic Microwave Background is highly valuable to constrain the growth of the structures in the Universe in a clean and robust manner over a wide range of redshifts. One of the theoretical systematics for lensing reconstruction is the impact of the lensing field non-Gaussianities on its estimators. Non-linear matter clustering and post-Born lensing corrections are known to bias standard quadratic estimators to some extent, most significantly so in temperature. In this work, we explore the impact of non-Gaussian deflections on Maximum a Posteriori lensing estimators, which, in contrast to quadratic estimators, are able to provide optimal measurements of the lensing field. We show that these naturally reduce the induced non-Gaussian bias and lead to unbiased cosmological constraints in Λ CDM at CMB-S4 noise levels without the need for explicit modelling. We also test the impact of assuming a non-Gaussian prior for the reconstruction; this mitigates the effect further slightly, but generally has little impact on the quality of the reconstruction. This shows that higher-order statistics of the lensing deflections are not expected to present a major challenge for optimal CMB lensing reconstruction in the foreseeable future.

I. INTRODUCTION

Gravitational lensing of the cosmic microwave background (CMB) is one of the leading cosmological probes of the next generation of CMB polarization surveys such as Simons Observatory (SO) [1] and CMB-S4 [2]. Maps of the CMB lensing potential and its summary statistics can provide clean and robust probes of the large-scale structure (LSS) in the Universe and allow to constrain cosmological parameters that govern the LSS growth and to which CMB anisotropies alone are weakly sensitive (for example neutrino mass, dark energy properties, curvature) [3–6]. On the other hand, galaxy surveys in different wavelengths probing the same LSS distribution (through galaxy clustering, weak lensing or intensity mapping) are also expected to extract complementary cosmological information on the LSS growth, or on primordial non-Gaussianities. Through galaxy-CMB lensing cross-correlations we will be able to marginalize over some observational systematics (e.g. shear multiplicative bias) [7–9], or modeling uncertainties (e.g. galaxy bias, magnification bias) [10–12]. Masses of high-redshift galaxy cluster samples can also be calibrated with the highest sensitivity through stacking of CMB lensing maps at location of the clusters. and later used to probe cosmology through their mass function [13–15]. Additionally, CMB lensing maps can also be used to predict and subtract the lensing-generated B -mode signal of CMB polarization to enhance constraints on inflationary physics achievable through precise measurements of the primordial B -mode signal on large angular scales [16, 17] Given the importance and diversity of the science case

connected to CMB lensing, it is crucial to understand all the properties and shortcomings of statistical estimators employed to reconstruct maps of the CMB lensing potential from the observed maps of the CMB anisotropies. The most commonly employed technique used for this purpose is the so-called quadratic estimator (QE) [18]. This uses the breaking of statistical isotropy introduced by the projected matter potential field into the observed CMB to reconstruct the lensing modes through weighted couplings between pairs of harmonic modes of the observed CMB itself. Such technique has been used to derive the most sensitive measurements in the field so far from the Planck satellite data or ground-based CMB polarization experiments such as ACTpol, SPTpol, Polarbear and BICEP [19–25].

Extensive effort has been recently carried out to evaluate the sensitivity to the QE to noise anisotropies, instrumental systematics, galactic and extragalactic foregrounds, and suitable modification have been proposed to minimize their impact in the final reconstructed lensing map and power spectrum for future experiments [25–31]. But for very deep surveys such as CMB-S4, the QE is expected to be suboptimal since it only accounts for the lensing coupling at linear order. Several additional methods that account for the full lensing information have been proposed. These advocate either sampling techniques [32–34], iterative maximum likelihood [35] or Maximum a-Posteriori (MAP) estimates [36] to reach the lowest possible reconstruction noise for future measurements of the lensing potential and of its power spectrum. MAP estimators in particular offer the advantage of a reduced computational cost compared to sampling or iterative spectrum reconstruction methods

and have so far been employed successfully on data covering large sky fraction. Very deep polarization observations of Polarbear were used to validate for the first time the performances of MAP estimates for joint lensing reconstruction and delensing analyses [37]. Similar results have then been achieved by sampling methods on the SPTpol data [38]. Recent work also showed how to achieve unbiased measurements of the CMB lensing spectrum from MAP estimates accounting for residual noise, normalization biases and mean field effects induced by noise anisotropies and masking, even in presence of mischaracterization of the data statistical properties [39, 40], so that it is now possible to start investigating biases induced by foregrounds and theoretical assumptions employed in the design of this new class of lensing estimator. In this work we focus our attention in particular on the assumption of the Gaussianity of the CMB lensing field for iterative MAP CMB lensing reconstruction. As shown in previous work of [41–43], non-Gaussian effects induced by the non-linear evolution of the LSS and post-Born lensing corrections due to multiple photon deflections [44] can bias the QE reconstruction ($N_L^{(3/2)}$ bias) and in turn affect the constraints achievable on cosmological parameters, in particular on the total mass of neutrinos [42]. Such bias becomes more important when correlating CMB lensing with external LSS tracers [45].

In Sec. II we review the properties of MAP estimator and $N_L^{(3/2)}$ for the QE. In Sec. IV we measure the $N_L^{(3/2)}$ bias in the MAP estimator, comparing it to the QE results and assess its importance for future surveys. In Sec. VIB we evaluate the impact $N_L^{(3/2)}$ on cosmological parameters and cross-correlation science and propose mitigation strategies.

In the following we use the standard convention and denote CMB lensing multipoles (L, M) and CMB multipoles with (ℓ, m) . We will also mention in the discussion the CMB lensing convergence κ , defined in function of the CMB lensing potential ϕ in real and harmonic domains as

$$\kappa(\hat{n}) = -\frac{\nabla^2}{2}\phi(\hat{n}) \quad (1)$$

$$\kappa_{LM} = \frac{1}{2}L(L+1)\phi_{LM} \quad (2)$$

II. LENSING RECONSTRUCTION AND NON-GAUSSIAN DEFLECTIONS

In this section, we provide background information on the lensing reconstruction methods we employ and offer a brief recap of the origin of the non-Gaussianity induced bias in the CMB lensing estimates.

First, we summarize the Maximum A Posteriori (MAP) CMB lensing estimator. For more details, we refer the reader to [46, 47] for the specific implementation we use.

A. Lensing reconstruction from the MAP estimator

To obtain the MAP estimator [46, 48, 49] we will start by modelling the observed CMB as

$$X^{\text{dat}} = \mathcal{B}\mathcal{D}_\alpha X + n, \quad (3)$$

where \mathcal{D}_α is the lensing operator mapping the primordial CMB X into the lensed CMB through the deflection α [48], \mathcal{B} a linear response matrix that includes the beam, and n the noise, that we assume to be uncorrelated with the CMB; in this work we will ignore foregrounds, and consider only experimental noise.

The goal of the MAP estimator is to derive a maximum a posterior estimate of the CMB lensing field, given the data, by maximising the posterior

$$p(\alpha|X^{\text{dat}}) \propto p(X^{\text{dat}}|\alpha)p_\alpha(\alpha), \quad (4)$$

given the likelihood $p(X^{\text{dat}}|\alpha)$ and a prior $p_\alpha(\alpha)$. For simplicity, due to the monotony of the natural logarithm, from now on will take the natural logarithm of probabilities, and so maximizing it will be equivalent to maximizing the original quantities.

We will assume a Gaussian log-likelihood for the data [46]

$$\ln p(X^{\text{dat}}|\alpha) = -\frac{1}{2}X^{\text{dat}} \cdot \text{Cov}_\alpha^{-1} X^{\text{dat}} - \frac{1}{2} \det \text{Cov}_\alpha + \text{const} \quad (5)$$

with a covariance

$$\text{Cov}_\alpha = \langle X^{\text{dat}} X^{\text{dat},\dagger} \rangle = \mathcal{B}\mathcal{D}_\alpha C^{\text{unl}} \mathcal{D}_\alpha^\dagger \mathcal{B}^\dagger + N, \quad (6)$$

where C^{unl} and N are the covariance of the unlensed CMB and that of the noise respectively, and we explicitly specify that the covariance has a structure that depends on the realization of the deflection angle α .

For now we will ignore the curl of the deflection field α , and we consider only its gradient mode. We can parametrize with κ , such that the final log-posterior for κ can be written as

$$\begin{aligned} -2 \ln p(\kappa|X^{\text{dat}}) &= X^{\text{dat}} \cdot \text{Cov}_\alpha^{-1} X^{\text{dat}} + \det \text{Cov}_\alpha \\ &+ \sum_{LM} \frac{\kappa_{LM} \kappa_{LM}^\dagger}{C_L^{\kappa\kappa}} \end{aligned} \quad (7)$$

up to irrelevant constants.

The CMB lensing estimator is then derived by requiring the gradient of the posterior, Eq. (7), to vanish. To solve for this, we will follow [46] and employ an iterative scheme starting from the standard quadratic estimator to finally obtain an estimate of the lensing potential field.

In particular, the calculation of the gradient of the likelihood gives a term that is quadratic in the data, which we will refer to as g^{QD} , and is given by the product of an inverse variance weighted field, and a *deflected* Wiener-filtered one [46]

$$g_{\text{QD}}(\hat{n}) = -\left(\mathcal{B}^\dagger \text{Cov}_\alpha^{-1} X^{\text{dat}}\right)(\hat{n}) \quad (8)$$

$$\cdot \left(\mathcal{D}_\alpha \delta C^{\text{unl}} \mathcal{D}_\alpha^\dagger \mathcal{B}^\dagger \text{Cov}_\alpha^{-1} X^{\text{dat}}\right)(\hat{n}),$$

where δ is the spin-raising operator.

The gradient of the likelihood also introduces a mean field term which, in the absence of other sources of anisotropies, represents the anisotropy introduced de-lensing the noise map by the estimated α in the quadratic gradient [46]. This term predominantly produces dilations (convergence-like) rather than local anisotropies (shear-like) terms, hence generally very small for polarization-only estimators [47]. However, it is larger for temperature estimators, possibly contributing up to 10–20% to the cross-spectrum with the input field. For further exploration of this impact, a brief discussion is presented in Appendix A, while a more detailed study is deferred to further work.

B. Lensing reconstruction from the QE estimator

The standard quadratic estimator (QE) [18] can be easily obtained from this likelihood perspective. If one writes the gradient (g_α) to linear order in α , and forces it to vanish, one obtains an estimate $\hat{\alpha}$ defined by:

$$g_\alpha \approx g_{\alpha=0}^{\text{QD}} - H_{\alpha=0} \hat{\alpha} = 0. \quad (9)$$

In this equation, H (the Hessian) is minus the second derivative of the log-likelihood function. Taking instead of the realization-dependent curvature its average of data realizations, $\hat{\alpha}$ will be quadratic in the data. This average is a Fisher matrix calculated at no deflection, and is identical to the quadratic estimator response function calculated in the standard manner [18].

In this case, the unnormalized estimate $g_{\alpha=0}^{\text{QD}}$ of the CMB lensing field is given by the product of an inverse-variance-weighted field, and an undeflected Wiener-filtered one

$$g^{\text{QE}}(\hat{n}) = \quad (10)$$

$$-\left(\mathcal{B}^\dagger \text{Cov}_{\alpha=0}^{-1} X^{\text{dat}}\right)(\hat{n}) \left(\delta C^{\text{unl}} \mathcal{B}^\dagger \text{Cov}_{\alpha=0}^{-1} X^{\text{dat}}\right)(\hat{n}). \quad (11)$$

This is similar to the quadratic part of Eq. (8), but for the absence of all deflections: the gradient part in (8) is deflected, because the iterative estimate works by capturing the residual lensing at the unlensed position, and then remapping back to give the estimate at the observed locations [50].

C. Non-Gaussian deflections effect on the reconstruction

The autospectrum of an estimated normalized QE CMB lensing map then results in

$$C_L^{\hat{\phi}\hat{\phi}} \approx C_L^{\phi\phi} + N_L^{(0)} + N_L^{(1)} + N_L^{(3/2)} + \dots, \quad (12)$$

where the total estimated power spectrum $C_L^{\hat{\phi}\hat{\phi}}$ has contributions from the signal of interest itself $C_L^{\phi\phi}$, from chance Gaussian fluctuations $N_L^{(0)}$ [3], from additional secondary lensing contractions $N_L^{(1)}$ [51], and non-Gaussian contributions $N_L^{(3/2)}$. These are induced by non-zero higher order statistics in the CMB lensing potential, arising from non-linear growth of structure and post-Born lensing [44].

The study of these effects on the estimate of the CMB lensing potential has been thoroughly studied for the quadratic estimator [41–43, 45]. They mainly arise due to a non-zero response in the auto-spectrum to the presence of a bispectrum term of the lensing potential.

For next generation surveys, such as CMB-S4 [2], the $N_L^{(3/2)}$ bias, if unaccounted for, can lead to 1-2 σ induced biases in cosmological parameters [42] such as the sum of neutrino masses, or have a large impact on cross-correlations with external matter tracers such as galaxy clustering or lensing [45].¹ While methods to mitigate these effects has been proposed, such as using polarization data only (by excluding temperature information that often leads to large biases), or alternative quadratic estimators, such as bias hardened ones [45, 52], these will not be able to fully harness the statistical power of upcoming CMB data as the MAP estimator does. Indeed, for low noise and high resolution experiments, the MAP will be able to reconstruct the lensing potential to the highest significance [39, 46], and therefore it is important to assess its potential in dealing with known lensing QE biases.

In this paper, we will focus on the quality of reconstruction of the MAP estimator, specifically examining its performance in the presence of non-Gaussian contributions. Our primary focus will be on the lensing auto-spectrum, which serves as the main observable for CMB lensing analyses. We additionally present results for the lensing cross-spectrum with the input lensing potential. We plan to investigate the effect on cross-correlations for MAP with large-scale structure surveys in future studies (as it was done for the QE case [45]).

Finally, we also present simulations results demonstrating the robustness of the bias hardening technique [52] to the non-Gaussian bias.

¹ See Section VIB for an updated discussion on the bias induced on the sum of neutrino masses.

III. MEASUREMENT SETUP

A. Experimental setup

We consider a crudely CMB-S4 wide-field-like experimental setup [2]. The assumed observed sky fraction is 40%, and the noise is modelled for simplicity as a homogeneous and isotropic white noise, with a noise level after component separation of N_{lev}^T (N_{lev}^P) = $1(\sqrt{2})\mu\text{K} - \text{arcmin}$ for the temperature T (polarization P), and a beam modelled as a Gaussian with full width at half-maximum of 1 arcmin. Our focus is to assess the fundamental performance of the MAP to the non-Gaussianity of CMB lensing, though it is worth noting that, at the small CMB scales considered here, extragalactic foregrounds will play an important role in temperature data.

These will still play a key role for future CMB lensing measurements like Simons Observatory [1], and are potentially the most sensitive reconstruction channel at even lower noise levels, provided scales deep enough in the damping tail of the CMB spectrum can be used, where the lensing effect is large. Figure 1 shows forecasts² for MAP estimators involving different data combinations in our baseline configuration, assuming CMB modes up to $\ell_{\text{max}} = 4000$ are used in the reconstruction. The polarization data on its own can reach a nominal precision of 0.25% on the spectrum using $L \leq 1500$ with the iterative approach, and provided all sources of foregrounds and complications are under control, one can push to slightly less than 0.2%.




FIG. 1. Predictions of the signal to noise ratio, defined in Equation 50, on the lensing auto spectrum for different lensing estimators, for our baseline CMB-S4 wide-like configuration, with $\ell_{\text{max}} = 4000$. Dashed is the standard QE and solid the MAP estimator. The temperature-only, polarization-only and minimum variance estimators are shown in red, orange and blue respectively.

² The MAP curves were calculated using 5 iterations, with theoretical reconstruction noise calculated using the `plancklens` code.

To extract the $N_L^{(3/2)}$ bias we will consider different kind of simulations, classified based on the nature of the input convergence field. To remove unphysical effects that could come from finite number of particles in the N-body simulation used in this work, we keep input lensing modes up to an $L_{\text{sim}} = 5120$. Given then an input CMB lensing map, we use `lenspyx` [53] to generate lensed CMB maps deflected with $\alpha = \nabla\phi$.³ Finally, we convolve the CMB simulation with the beam, and add a noise realization. The total map $X^{\text{obs}} = X^{\text{cmb}} + n$ will be our observed map.

We now proceed to describe the sets of deflection fields used to lens the CMB.

B. Simulated data sets

In order to isolate the relevant effects and test the impact of the assumption of the prior on κ for the lensing reconstruction, we made use of different sets of simulations, summarized below:

- $\kappa^{\text{tot}} \equiv \kappa^{\text{LSS+PB}}$, CMB lensing convergence simulations that include the total effect from LSS non-linearity and post-Born (PB) effects.
- κ^{LSS} , CMB lensing convergence simulations that include only the effect from LSS non-linearity.
- $\kappa^{X,\text{R}}$, with $X \in \{\text{tot}, \text{LSS}\}$ where we randomize the phases of the CMB lensing convergence simulations that include the total effect from LSS non-linearity and PB effects, or only LSS if $X = \text{LSS}$.
- κ^{G} , Gaussian CMB lensing simulations that include the exact power spectrum of the raw input convergence field $C_L^{\kappa^X \kappa^X}$, $X \in \{\text{tot}, \text{LSS}\}$.⁴

In addition, in Section III B 2 we will also use log-normal simulations of κ with a custom third moment γ , and they will be indicated by κ^{log} , and their randomized version by $\kappa^{\text{log,R}}$.

Below we provide more details on each type of simulations.

1. Fully non-linear observables

To model the realistic effect of the non-linear LSS clustering of matter and the effect of post-Born corrections on

³ <https://github.com/carronj/lenspyx>

⁴ Though we generate Gaussian simulations with CMB lensing power spectra matching the total and LSS only cases, we note that the impact of post-Born corrections on $C_L^{\kappa\kappa}$ has been proven to be negligible at the level of accuracy considered in this work, at the 0.25% level at $L = 8000$ [44, 54]. Therefore one might only use only Gaussian simulations with LSS only spectra.

lensing observables we used the full-sky maps of the CMB lensing and curl potential of [54]. These were constructed using a multiple-lens plane raytracing algorithm [55] with lensing planes constructed from a Λ CDM simulation of the DEMNUni suite. This was designed to study the impact of massive neutrinos on the universe evolution and its interplay with different dark energy models [56, 57]. The Λ CDM simulation used a Planck 2013 cosmology with massless neutrinos

$$\{\Omega_{\text{cdm}}, \Omega_{\text{b}}, \Omega_{\Lambda}, n_{\text{s}}, \sigma_8, H_0, M_{\nu}, \tau\} = \quad (13)$$

$$\{0.27, 0.05, 0.68, 0.96, 0.83, 67 \text{ Km/s/Mpc}, 0, 0.0925\},$$

and sampled the matter distribution with 2048³ dark matter particles in a volume of 2Gpc/h between $z = 99$ and $z = 0$. The mass resolution of the simulation at $z = 0$ is $M_{\text{CDM}} = 8.27 \times 10^{10} M_{\odot}/h$. The details of the full-sky lightcone construction from the finite volume of the N-body simulation are provided in Refs. [58, 59]. The output of the lightcone construction consists in 62 surface mass density planes $\Sigma^{\theta(k)}$ including the mass contained in spherical shells of comoving thickness $\Delta\chi \approx 150 \text{ Mpc}/h$ that are used to construct k-th CMB lensing potential convergence plane as

$$\Delta_{\Sigma}^{(k)} = \Sigma^{\theta(k)} / \bar{\Sigma}^{\theta(k)} - 1. \quad (14)$$

$$\kappa_{\chi_{\text{CMB}}}^{(k)} = 4\pi G \frac{D_a(\chi_{\text{CMB}} - \chi_k)}{D_a(\chi_{\text{CMB}})} \frac{(1 + z_k)}{D_a(\chi_k)} \Delta_{\Sigma}^{(k)} \quad (15)$$

where D_a is the angular diameter distance and $\chi_{\text{CMB}}, \chi_k$ are comoving distances to the CMB and to the k -th lensing plane respectively. The simulation neglects the matter distribution at $99 < z \leq 1089$ without any effective loss in accuracy. The lensing convergence planes can then be summed together to obtain the lensing potential in Born approximation (κ^{LSS}) following [60] or used to propagate the full lensing distortion tensor beyond the Born approximation as discussed in [54] and obtain also maps of CMB κ^{tot} and lensing rotation ω which arises from coupling between subsequent lensing events. The fact that both Born and post-Born lensing maps are derived from the exact same matter distribution allows to disentangle the impact of each of the specific term as anything depending on $\kappa^{\text{tot}} - \kappa^{LSS}$ will isolate the effect of post-Born corrections alone. The pipeline is general and can be adopted for lensing planes located at χ_s different from the CMB. We refer the reader to [54, 61] for more technical details of the raytracing procedure used here.

2. Log-normal simulations

The main relevant effect for our study is the bispectrum signal introduced in the N-body derived convergence map. For comparison, it will be useful to generate simpler maps that introduce a connected higher than two order correlation function, such as log-normal simulations. These are much cheaper to produce and better

understood analytically compared to full N-body simulations, while allowing to tune the skewness of the generated map to match the value found in the full N-body simulation.

We follow the methods of [62] to produce our log-normal simulations.⁵ We model the CMB lensing with a shifted lognormal field:

$$\kappa^{\text{log}}(\hat{n}) = e^{Z(\hat{n})} - \lambda \quad (16)$$

where Z is a Gaussian random field with mean μ and variance σ^2 , and λ a shift parameter. These values are calculated in such a way to match a desired power spectrum and skewness calculated from a band-limited map. We relate the moments of the desired map $\mu_{\kappa^{\text{log}}}$, $\sigma_{\kappa^{\text{log}}}^2$, and $\gamma_{\kappa^{\text{log}}}$ to the parameters μ , σ , and λ to generate the log-normal simulation (see e.g. [62])

To obtain the desired log-normal field, we first calculate the correlation function of the convergence field, and obtain from it the correlation function of the Gaussian field Z according to [62]

$$\xi_Z = \ln \left(\frac{\xi_{\kappa^{\text{log}}}}{\alpha^2} + 1 \right) \quad (17)$$

where $\alpha = \mu_{\kappa^{\text{log}}} + \lambda > 0$. We then generate the zero mean Gaussian field $Z - \mu$, from the power spectrum C_l^Z , obtained from ξ_Z ,⁶ using `Healpy` [63, 64].⁷ Finally, we go back to real space and obtain

$$\kappa_{\text{log}}(\hat{n}) = e^{Z(\hat{n})} - \lambda, \quad (18)$$

where $e^{\mu} = (\mu_{\kappa^{\text{log}}} + \lambda)e^{-\sigma^2/2}$, with $\sigma^2 = \xi_Z(0)$ is the variance of the Gaussian field Z . The generated maps are saved at the same resolution NSIDE = 4096 of the N-body convergence simulation, and are then processed with the same pipeline. In Figure 2 we show on the top panel an histogram of the non-Gaussian input and log-normal maps, and on the lower panel the third moment for the different simulations used in this work. As we will discuss in detail later in the text, the bulk of non-Gaussian biases is related, to leading order, to the bispectrum of the lensing field and recent work has shown that shifted-lognormal simulations like the ones adopted here can provide a moderately accurate approximation to weak lensing fields' higher-order statistics [65].

⁵ As implemented in <https://github.com/Saladino93/fieldgen>.

⁶ We use the package <https://cltools.readthedocs.io/flt/> to perform Legendre transforms to quickly switch from/to correlation functions/angular power spectra.

⁷ healpy.readthedocs.io/




FIG. 2. *Upper panel:* binned histogram of several input convergence maps. We can see that the histograms from the LSS and lognormal convergence maps are slightly skewed compared to the Gaussian and randomized ones. We do not plot an histogram of a convergence map including LSS+Post Born effects, though this will have a reduced skewness compared to the LSS only one [54]. *Lower panel:* The skewness for different input convergence maps in function of the FWHM of a Gaussian smoothing. The lognormal case shown here is built to match the log-normal parameter λ of the LSS map at $\text{fwhm} = 0$. The values are calculated considering lensing modes up to $L_{\text{sim}} = 5120$ where the CMB lensing modes used to lens the CMB simulations are cut at.

C. Lensing potential reconstruction setup

Unless indicated otherwise, we reconstruct the lensing potentials for $L \in (2, 5120)$ from the generated simulations, using CMB modes in the range $\ell_{\text{T,CMB}} \in (10, 4000)$ for T , E and B .

We use `delensalot` [47]⁸ for the MAP or QE reconstruction of the lensing potentials (the QE reconstruction of `delensalot` follows `plancklens`). Unless stated otherwise, throughout the paper the presented MAP results performed at least 5 iterations, where we find converged

results, in the sense of negligible changes in the derived spectra of the reconstructed fields.⁹

For the MAP estimates, reconstruction of lensing potential and the Wiener filtered CMBs are attempted up to an $\ell_{\text{unl,max}} = 5120$ at each step. The step size used in the Newton-Raphson method to updated the iterative estimates is $\lambda = 0.5$, and the tolerance to check for convergence of the conjugate gradient method used for the inverse variance operation is set to $\text{cg_tol} = 10^{-7}$, running 7-10 iterations per reconstruction.

We renormalize the reconstructed maps as follows: we do not employ analytical expressions, but use reconstructions based on observed CMB simulations lensed with the Gaussian CMB lensing simulations set (κ^{G}), and we calculate a normalization factor given by

$$\mathcal{W}_L^{\text{emp,XY}} = \frac{\langle C_L^{\hat{g}^{XY} \kappa} \rangle_{\text{sims}}}{\langle C_L^{\kappa \kappa} \rangle_{\text{sims}}},$$

where \hat{g}^{XY} represents an unnormalized CMB lensing potential reconstruction, and κ the corresponding input used to lens the CMB simulation.¹⁰

We then use this isotropic factor to normalize the lensing potential estimates from all simulation sets ($\kappa^X, \kappa^{X,R} \mid X \in \{\text{LSS, T, log}\}$)

$$\hat{\kappa}_{LM}^{XY} = \frac{\hat{g}_{LM}^{XY}}{\mathcal{W}_L^{\text{emp,XY}}} \quad (19)$$

from which we calculate the raw power spectrum

$$C_L^{\hat{\kappa}^{XY} \hat{\kappa}^{XY}} = \frac{1}{2L+1} \sum_M \hat{\kappa}_{LM} \hat{\kappa}_{LM}^\dagger. \quad (20)$$

Finally, to assess the quality of the reconstructions, it is useful to define the cross correlation coefficient ρ_L ¹¹.

$$\rho_L = \frac{C_L^{\hat{\kappa}^{XY} \kappa}}{\sqrt{C_L^{\hat{\kappa}^{XY} \hat{\kappa}^{XY}} C_L^{\kappa \kappa}}},$$

⁹ In the case of temperature, convergence of the results can be reached earlier though we still keep a larger number of iterations.

¹⁰ We checked that for the QE, we get consistent results both with this method and by using a response function given by CMB gradient spectra, following [45].

¹¹ By ‘quality’ of the reconstruction we mean here its information content on the true lensing potential, irrespective of potential biases (the cross-correlation coefficient is invariant under arbitrary multiplicative biases in the reconstructed field (small additive ones can also be cast as multiplicative biases))

⁸ <https://github.com/NextGenCMB/delensalot>

IV. $N_L^{(3/2)}$ BIAS

With the notation now established, we are ready to calculate the biases associated with the presence of non-Gaussianity in our CMB lensing fields. The $N_L^{(3/2)}$ biases are determined through simulations, specifically by cross-correlating the estimated CMB lensing field with another tracer, such as the input CMB lensing potential itself, or by examining the CMB lensing auto-correlation. This methodology, as outlined in ([42, 45]) allows us to quantitatively assess and account for the impact of non-Gaussian features in our lensing fields. This is achieved by calculating the difference from the scenario where no non-Gaussianity is present.

The biases, denoted as $N_L^{(3/2)}$, are computed through the following expressions:

$$\begin{aligned} \text{Total : } N_L^{(3/2)} &= \langle \hat{C}_L^{\phi^{XY} \phi^{\text{ext}}}[\kappa^{\text{tot}}] - \hat{C}_L^{\phi^{XY} \phi^{\text{ext}}}[\kappa^{\text{tot,R}}] \rangle_{\text{CMB}} \\ \text{LSS : } N_L^{(3/2)} &= \langle \hat{C}_L^{\phi^{XY} \phi^{\text{ext}}}[\kappa^{\text{LSS}}] - \hat{C}_L^{\phi^{XY} \phi^{\text{ext}}}[\kappa^{\text{LSS,R}}] \rangle_{\text{CMB}} \\ \text{PB : } N_L^{(3/2)} &= \langle \hat{C}_L^{\phi^{XY} \phi^{\text{ext}}}[\kappa^{\text{tot}}] - \hat{C}_L^{\phi^{XY} \phi^{\text{ext}}}[\kappa^{\text{LSS}}] \rangle_{\text{CMB}} . \end{aligned} \quad (21)$$

We denote ϕ^{ext} either the input lensing field or the reconstructed field $\hat{\phi}^{XY}$ itself. The angle brackets signify the average over lensed CMB simulations. To gauge the reliability of our measurements, we will compute the scatter in the biases from these simulations to obtain the uncertainty on the mean measurement. Our subtraction method ensures that simulations which are differentiated in these equations share identical primordial CMB and noise, mitigating realization-dependent biases and cosmic variance effects.

In Figures 3 and 4 we show the $N_L^{(3/2)}$ bias for the CMB lensing autospectrum and cross-spectrum with the input, respectively. The shaded areas are computed from the scatters on the mean of the simulations, while grey area represent statistical error bars as calculated from the diagonal of the covariance matrix

$$C_{LL'} = \frac{2\delta_{LL'}}{(2L+1)f_{\text{sky}}} \left(C_L^{\phi\phi} + N_L^{(0)} + N_L^{(1)} \right)^2 . \quad (22)$$

In Figure 3, it is evident that, for the lensing autospectrum, the $N_L^{(3/2)}$ bias of the MAP estimator shows a lower absolute value compared to the QE estimator for $L \leq 1800$. Specifically, at the scales presented, the LSS

effect in the QE tends to suppress power in the reconstructed potential. This effect is alleviated in the MAP for all estimators except on the smallest scales. The situation is similar, but with the opposite sign, for the PB bispectrum effect. This results in partial cancellation between the two effects for the total contribution for both the QE and MAP. Similar conclusions are achieved with the cross-spectrum with the input lensing potential as depicted in Figure 4.

Both MAP and QE estimators share the characteristic that the total bias effect depends on the combination of used data. Specifically, in the minimum variance combination, the polarization estimator dominates at lower modes $L \leq 1500$, while the temperature one dominates at higher L 's.

Finally, a distinct feature is the noticeable steep rise in the bias of the MAP estimator for the LSS cases, particularly in the TT and MV combinations, evident in both the auto and cross spectra. While the specific origin of this rise remains uncertain, we validated our bias method by assessing Gaussian simulations with no LSS bias and randomized ones. We checked that the difference is consistent with zero. Therefore this rise at small scales in the MAP might be due non-Gaussianities, or from the choice of normalisation of the maps (obtained from Gaussian simulations). To investigate the choice of normalisation, we consider the cross-correlation coefficient of the reconstructions with their respective inputs

$$\rho^X = \frac{\langle C_L^{\hat{\kappa}\kappa} \rangle}{\sqrt{\langle C_L^{\hat{\kappa}\hat{\kappa}} \rangle \langle C_L^{\kappa\kappa} \rangle}} \quad (23)$$

where $\langle \cdot \rangle$ denotes the average over sims. This quantity does not require any normalisation, and it is related to the faithfulness of our reconstruction to the input. We assume that we can directly take the difference of the cross-correlation coefficient with the corresponding input for the LSS non-Gaussian MAP vs the Gaussian MAP estimates, as shown in Figure 5.

As expected, we can see that on large scales the Gaussian case reconstruction is better correlated with the input, for both the QE and MAP cross-correlation differences. While on small scales, $L > 1500$, interestingly we see a small rise in the difference with the cross-correlation coefficient of the MAP of the non-Gaussian case. This might be related to additive terms that we test in Section IV B 2.

A. Simple theory calculations

Given the complexity of beyond-the-QE reconstruction, a comprehensive analytic treatment of the theory

curves for the MAP- $N_L^{(3/2)}$ bias seems out of reach. Nevertheless, we can proceed by analogy with quadratic estimator theory to obtain some curves that matches reasonably well our findings.

We will focus on the simpler case of the cross-




FIG. 3. Fractional $N_L^{(3/2)}$ bias on the CMB lensing auto-correlation for a CMB-S4 like configuration, as described in the text. This is calculated by averaging over 64 realization of the input CMB primordial field. We show results for the LSS only, post-Born (PB) and full cases in blue, green and red, respectively. In grey and light grey we show the statistical error bars for MAP and QE respectively. We bin the spectra from $L_{\min} = 30$ to $L_{\max} = 3000$ with a wide binning of around 140.




FIG. 4. Same as Figure 3 but for the cross-spectrum between the reconstruction and the input lensing potential.

correlation of the reconstructed CMB lensing potential $\hat{\phi}$ with the input lensing potential ϕ on the flat-sky:¹²

$$N_L^{(3/2),\text{cross}} = \langle \hat{\phi}(\vec{L})\phi^*(\vec{L}) \rangle - \langle \phi(\vec{L})\phi^*(\vec{L}) \rangle, \quad (24)$$

where we assume the reconstruction $\hat{\phi}(\vec{L})$ is unbiased, in the sense that $\langle \hat{\phi} \rangle = \phi$. The difference between these two

¹² We focus on the cross-correlation with the input, and not on the reconstruction auto-spectrum, as the former is easier to predict with respect to the latter. In particular, we perform all of our studies using CMB modes up to $l_{\max} = 4000$. Our baseline perturbative analytical model for predicting the QE auto-spectrum, based on [41, 45], works well only on the largest scales, where we get the prediction $N^{(3/2),\text{auto}} \sim 2N^{(3/2),\text{cross}}$. This will require a more careful treatment for future studies.

terms depends predominantly on the bispectrum of the observed lensing potential $B^{\phi\phi\phi}(L, l_2, l_3)$, where L stands for the external tracer (here ϕ itself) multipole and l 's for the ones entering through the CMB modes used for CMB lensing reconstruction.

For the standard QE CMB lensing temperature-only estimator, calculations are given in [41, 45]. We briefly review the rationale here.

A temperature-based quadratic estimator is of the form

$$\hat{\phi}(\vec{L}) = A_L^{TT} \sum_{\vec{l}} g(\vec{l}, \vec{L}) T^{\text{dat}}(\vec{L} - \vec{l}) T^{\text{dat}}(\vec{l}). \quad (25)$$

for some normalization A_L^{TT} and weights $g(\vec{l}, \vec{L})$.




FIG. 5. Difference in the cross-correlation coefficients between the LSS non-Gaussian (blue), randomized LSS non-Gaussian (orange) cases with the Gaussian case, for the TT estimator. In solid (dashed) we have the MAP (QE). We can see that the difference with randomized (orange) is consistent with zero. While as expected, the non-Gaussian case is below zero on large scales. Nevertheless, for the non-Gaussian MAP (blue solid line), we can see a small rise after $L > 1500$.

The lensed CMB may be perturbatively expanded in a series with respect to the lensing potential, leading to

$$T = T^u + \delta T^u + \delta^2 T^u + \mathcal{O}(\phi^3), \quad (26)$$

where T^u is the unlensed CMB, and $\delta^n T^u$ depends on the lensing potential power n [3].

On evaluating the cross-correlation (24),

$$\langle \hat{\phi}(\vec{L})\phi(-\vec{L}) \rangle \sim \langle T(\vec{L} - \vec{l})T(\vec{l})\phi(-\vec{L}) \rangle, \quad (27)$$

the lensing potential bispectrum will give rise to the $N_L^{(3/2),\text{cross}}$,

$$N^{(3/2),\text{cross}} \sim \langle \delta T^u \delta T^u \phi \rangle + 2 \langle T^u \delta^2 T^u \phi \rangle. \quad (28)$$

The full expression for the temperature case, including to a good approximation non-perturbative lensing remapping effects [45], is

$$\begin{aligned} N_L^{(3/2),\text{cross}} = & A_L^{TT} \int_{\vec{l}_1} B^{\phi\phi\phi}(L, l_1, |\vec{L} - \vec{l}_1|) \int_{\vec{l}_2} g(\vec{l}_2, \vec{L}) \\ & \left(-C_{|\vec{l}_1 - \vec{l}_2|}^{T\nabla T} [(\vec{l}_1 - \vec{l}_2) \cdot \vec{l}_1][(\vec{l}_1 - \vec{l}_2) \cdot (\vec{L} - \vec{l}_1)] \right. \\ & \left. + C_{\vec{l}_2}^{T\nabla T} [\vec{l}_2 \cdot \vec{l}_1][\vec{l}_2 \cdot (\vec{L} - \vec{l}_1)] \right). \end{aligned} \quad (29)$$

The spectra $C_l^{T\nabla T}$ are here the same lensed gradient spectra that enters the non-perturbative lensing response functions.

Let's discuss how we connect the MAP bias to this representation. Reference [39] demonstrated through simulations that the converged CMB lensing MAP solution

power spectrum (for polarization at least) can be accurately described as a quadratic estimator with partially lensed CMB spectra. These spectra can be obtained using an iterative scheme initially proposed for the EB estimator by [17]. Hence, it is natural to test this recipe for $N_L^{3/2}$ as well.

In this picture, the lensing potential entering the CMB legs are now given by the unresolved, residual lensing map

$$\hat{\phi}_{LM}^{\text{res}} \equiv (1 - \mathcal{W}_L)\phi_{LM}. \quad (30)$$

We then make use of the same flat-sky prediction Eq (29), but with the substitution

$$B^{\phi\phi\phi}(L, l_1, l_3) \rightarrow B^{\phi\phi\phi}(L, l_1, l_3)(1 - \mathcal{W}_{l_1})(1 - \mathcal{W}_{l_3}),$$

where $\vec{l}_3 = \vec{L} - \vec{l}_1$, and \mathcal{W} is the Wiener-filter. The normalization A_L^{TT} is the standard quadratic estimator normalization but calculated with partially lensed gradient spectra, and $C_l^{T\nabla T}$ are also calculated with the partially lensed spectra.

In Figure 6 we show the result of this naive calculation, compared to simulations.¹³ For simplicity we consider the case of temperature only. On the upper panel we show the LSS-only case, where $N^{3/2}$ is stronger than with the total bispectrum, shown in the lower panel. It assumes a LSS bispectrum calculated at the tree level in perturbation theory with corrections coming from [66] (though an improved version can be found in [67]). The Post-Born corrections are based on [44]. See the Appendix of [45] for a short review.

We can see that on a wide range of scales this simple predictive scheme is able to recover the estimated biases from simulations, validating the understanding that the bispectrum of the CMB lensing field correctly describe the bias also in the case of the iterative estimator. In the LSS-only case, and on small scales, the prediction deviates significantly from our findings in simulations. In Section IV B 2 we check that its origin is not from higher even CMB lensing connected n -point functions.

B. Consistency and robustness tests

To have a better understanding of our reconstructions and validate our pipeline, we perform a series of consistency checks to explore possible residuals in our estimates, specifically those of higher order than the bispectrum residuals.

1. Joint reconstruction with lensing field rotation

In this section we test the impact of post-Born lensing rotation on the MAP reconstructions.

¹³ The code to calculate these biases is based on <https://github.com/Saladino93/lensbiases>.




FIG. 6. Approximate theory calculation for the $N_L^{(3/2)}$ bias on the CMB lensing cross-spectrum. In blue, we show results for the quadratic estimator, and in red for the MAP reconstruction. The MAP predictions are obtained from the naive analytic prescription given in the text. This is for our CMB-S4-like reconstruction from temperature.

The CMB lensing deflection vector field $\vec{\alpha}(\hat{n})$ can be written thanks to the Helmholtz decomposition in flat-sky notation as

$$\vec{\alpha}(\hat{n}) = \vec{\nabla}\phi(\hat{n}) + \vec{\nabla} \times \Omega(\hat{n}), \quad (31)$$

On the full-sky, the decomposition is compactly written using the spin-weight formalism,

$${}_1\alpha(\hat{n}) = -\bar{\partial}\phi(\hat{n}) - i\bar{\partial}\Omega(\hat{n}), \quad (32)$$

where Ω is the lensing curl potential, and

$$\begin{aligned} \omega(\hat{n}) &= -\frac{\nabla^2\Omega(\hat{n})}{2} \\ \omega_{LM} &= \frac{1}{2}L(L+1)\Omega_{LM} \end{aligned} \quad (33)$$

is the lensing field rotation; the angle by which tiny local images are rotated by lensing. The leading rotation is induced by post-Born lensing, that couples pair of non-aligned shearing lenses at different redshifts, producing in this way a net rotation. For a single deflection or in the Born approximation the rotation is null. Non-zero bispectra of the kind $\kappa\kappa\omega$ or $\kappa\omega\omega$ [44, 54] are also generate at higher perturbative orders. The amplitude of the rotation field power spectrum is about 3 to 4 orders of magnitude smaller than the one of κ [35, 44] on the scales considered here.

The simulations of [45, 54] naturally include the rotation component of CMB lensing. We use this field in combination with the post-Born convergence map to estimate the lensing potentials in the presence of a curl-like displacement using either a joint ϕ - Ω optimal reconstruction (the differences to the standard ϕ -only algorithm are minors and described in [47]), or completely neglecting Ω in the reconstruction.

We can then estimate the effects in the $N_L^{(3/2)}$ bias coming from the mixed-bispectrum terms $\kappa\kappa\omega$, $\kappa\omega\omega$ as :

$$\text{Mixed : } N_L^{(3/2)} = \langle \hat{C}_L^{\phi^{XY}} \phi^{\text{ext}}[\kappa^{\text{tot}}, \omega] - \hat{C}_L^{\phi^{XY}} \phi^{\text{ext}}[\kappa^{\text{tot}}] \rangle_{\text{CMB}}, \quad (34)$$

where the first term include the lensing rotation in the inputs and reconstruction, and the second does not. As for the biases obtained in previous sections, here we use 64 CMB primordial realizations, and we run the pipeline up to 10 iterations, though outputs appeared to have converged starting from the fifth one.

We show in Fig. 7 the $N_L^{(3/2)}$ biases to the auto and cross-spectra in the presence of a full deflection field (including LSS and post-Born effects in the gradient-like deflection component as well as curl-like deflections), compared to a case without curl-like displacement component in the simulations. We see no statistically significant differences.




FIG. 7. Non-Gaussian $N_L^{(3/2)}$ lensing biases in the estimated lensing potential auto-spectra (top panels) or cross-spectrum to the true lensing (bottom panels), with and without curl-like deflections in the simulations. In red we show the total bias that does not include post-Born lensing field rotation (same as Figure 3), while in orange we show the difference between the biases obtained including post-Born corrections for both the gradient and curl component of the deflections and the red curve. The difference is consistent with zero in both panels.

2. Sign-flipped κ

At the perturbative level, the $N_L^{(3/2)}$ bias is proportional to a projected bispectrum on large scales [41]. We will test this assumption in this section.

Flipping the sign of the lensing potential should flip the sign of the bias, if it depends the bispectrum that is odd in the lensing field. For the QE, this argument breaks down only when using the smallest CMB scales for our

reconstructions, where perturbative arguments are less effective (e.g., at $\ell_{\max} \sim 4000$ [41]). On the other hand, the MAP always uses the full likelihood information and, therefore, higher-order point functions to reconstruct the lensing field. This also makes the MAP more difficult to assess analytically compared to the QE. Therefore, we turn to simulations to check if the MAP bias behaves significantly differently with respect to the sign of the input lensing map.

For each of the convergence maps (κ^G , κ^{LSS} , κ^{tot}) we calculate their flipped version

$$\kappa_{LM}^{\text{in},\neg} = (-1) \cdot \kappa_{LM}^{\text{in}}, \quad (35)$$

where κ_{LM}^{in} is the reference input CMB convergence lensing field. We generate then 64 lensed CMB simulations with this flipped lensing potential, and calculate the half-sum and half-differences

$$\Delta_{\pm}^{XY} = \left\langle \frac{1}{2} \left(N_L^{(3/2)}[\kappa_{LM}^{\text{in}}] \pm N_L^{(3/2)}[\kappa_{LM}^{\text{in},\neg}] \right) \right\rangle_{\text{sims}}. \quad (36)$$

for the TT , EE , $EE + EB$, and MV estimators, and where the $N_L^{(3/2)}$ are estimated as explained in Section IV.

Any contribution from contractions which are even in the lensing field will appear in the half-sum of the biases, and similarly for odd terms in the half-differences. Figure 8 shows these half-sum in blue and half-differences in orange, for the cross LSS-only reconstruction bias from temperature. We can see that for MAP and QE reconstructions we have a small residual in the half-sum ($\sim 0.2\%$). In particular, the rise at the highest L observed in the previous section is still most likely to an odd n -point function effect.

To quantify the consistency of these results with a signal dominated by the bispectrum, we calculate the probability-to-exceed (PTE) using Welch's test. This test compares the means of two simulation sets (κ , sign-flipped κ) with unequal variances, calculated from simulation scatters. The input data vector is the average over simulation measurements for 21 multipole bins over the range $5 \leq L \leq 3000$. The results are shown in Table I, where the average PTEs over the 21 multipoles are shown. A test is assumed successful if the PTE is greater than 0.05. With this definition, the QE MV , EB cross-spectrum for PB cases fail, though this is acceptable given the number of tests performed.

V. ALTERNATIVE CMB LENSING ESTIMATORS

A. Iterative estimator with a non-Gaussian prior

The conventional iterative estimator for CMB lensing assumes a Gaussian potential field due to its incorporation in the prior distribution. However, the observed




FIG. 8. Half-sum and half-differences of $N_L^{(3/2)}$ biases, built from a baseline simulation, and a second one flipping the sign of the input lensing potential, in order to isolate the contributions of even (blue) and odd (orange) n -point functions on the total bias. The light (dark) colors are the QE (MAP) reconstructions for $\ell_{\max} = 4000$, from temperature-only and LSS bispectrum only.

CMB lensing field exhibits non-Gaussian characteristics that the standard iterative approach does not directly address. In this context, our objective is to accommodate potential non-Gaussian aspects of the field by proposing an alternative method for estimating ϕ .

To begin with, we introduce a non-Gaussian prior. Specifically, we characterize the lensing convergence field using a log-normal model, as outlined in section III B 2. This choice establishes the foundation for our new approach and is motivated by the fact that the log-normal approximation is a sensible one for moderately non-linear fields, e.g. [68], [69].

The basic idea is that instead of iteratively solving for the CMB lensing potential field directly, we iterate over a posterior that is a function of the Gaussian field whose exponential gives the CMB lensing convergence field.¹⁴

In this case, we will assume a prior on the Gaussianized field Z , which may be written in harmonic space

$$-2 \ln p_Z(Z) = \frac{(Z_{00}^2 - \mu\sqrt{4\pi})^2}{C_{L=0}^Z} + \sum_{L \geq 0, |M| \leq L} \frac{Z_{LM} Z_{LM}^\dagger}{C_L^Z}, \quad (37)$$

up to irrelevant constants. The un-normalized posterior can be written as before

$$p(Z|X^{\text{dat}}) \propto p(X^{\text{dat}}|Z)p(Z) \quad (38)$$

¹⁴ This allows us to use our baseline MAP pipeline with little modifications.

Estimator	QE		MAP	
	auto	cross	auto	cross
	LSS / TOT / PB	LSS / TOT / PB	LSS / TOT / PB	LSS / TOT / PB
	$\ell_{\max} = 3000$			
<i>TT</i>	0.51 / 0.78 / 0.31	0.56 / 0.21 / 0.24	0.89 / 0.51 / 0.98	0.66 / 0.20 / 0.33
<i>EE</i>	0.42 / 0.28 / 0.95	0.44 / 0.63 / 0.56	0.32 / 0.19 / 0.86	0.65 / 0.79 / 0.72
<i>EE + EB</i>	0.11 / 0.17 / 0.69	0.07 / 0.61 / 0.04	0.58 / 0.96 / 0.25	0.10 / 0.28 / 0.10
<i>MV</i>	0.90 / 0.47 / 0.13	0.11 / 0.79 / 0.02	0.20 / 0.10 / 0.94	0.75 / 0.74 / 0.45
	$\ell_{\max} = 4000$			
<i>TT</i>	0.97 / 0.87 / 0.88	0.74 / 0.71 / 0.75	0.38 / 0.43 / 0.37	0.25 / 0.09 / 0.36
<i>EE</i>	0.07 / 0.06 / 0.43	0.31 / 0.28 / 0.37	0.52 / 0.29 / 0.74	0.28 / 0.27 / 0.35
<i>EE + EB</i>	0.38 / 0.99 / 0.31	0.15 / 0.44 / 0.14	0.82 / 0.88 / 0.65	0.32 / 0.59 / 0.22
<i>MV</i>	0.42 / 0.96 / 0.14	0.25 / 0.25 / 0.27	0.99 / 0.93 / 0.95	0.77 / 0.73 / 0.79

TABLE I. Averaged over L -bins PTEs for the half-sum tests of the $N_L^{(3/2)}$ biases. Shown are the temperature, polarization, their combination, QE and MAP estimators, for $N_L^{(3/2)}$ auto- and cross spectra.

Subsequently, our objective is to determine the optimal estimate of Z based on the available data.¹⁵

In this case, the total gradient with respect to the Gaussian field is

$$g_Z^{\text{tot}}(\hat{n}) = \frac{\delta \ln p(Z|X^{\text{dat}})}{\delta Z(\hat{n})} = g_Z^{\text{QD}}(\hat{n}) - g_Z^{\text{MF}}(\hat{n}) + g_Z^{\text{PR}}(\hat{n}), \quad (39)$$

This gradient can be calculated with minimal modifications to our ϕ -based MAP reconstruction code: using the chain rule, we have namely

$$\begin{aligned} g_Z^{\text{QD}}(\hat{n}) - g_Z^{\text{MF}}(\hat{n}) &= \frac{\delta \ln p(X^{\text{dat}}|Z)}{\delta Z(\hat{n})} \\ &= \int d^2 \hat{n}' \frac{\delta \kappa(\hat{n}')}{\delta Z(\hat{n})} \frac{\delta \ln p(X^{\text{dat}}|\phi)}{\delta \kappa(\hat{n}')} \\ &= e^{Z(\hat{n})} (g_{\kappa}^{\text{QD}}(\hat{n}) - g_{\kappa}^{\text{MF}}(\hat{n})) \end{aligned} \quad (40)$$

We have used in the second line

$$\frac{\delta \kappa(\hat{n}')}{\delta Z(\hat{n})} = \delta_D^{(2)}(\hat{n}' - \hat{n}) e^{Z(\hat{n})}, \quad (41)$$

which follows directly from our definition $\kappa(\hat{n}) = e^{Z(\hat{n})} - \lambda$. Finally, the gradients $g_{\kappa}(\hat{n})$ are easily obtained in harmonic space from those of ϕ which the original MAP reconstruction calculates: from $\kappa_{LM} = \frac{1}{2}L(L+1)\phi_{LM}$ follows for the gradients

$$g_{\kappa}(\hat{n}) = \int d^2 n' \frac{\delta \phi(\hat{n}')}{\delta \kappa(\hat{n})} g_{\phi}(\hat{n}') \quad (42)$$

$$= \sum_{LM} \frac{2}{L(L+1)} g_{\phi,LM} Y_{LM}(\hat{n}). \quad (43)$$

Starting from the quadratic estimator solution we then iterate over estimates of Z in the same manner than the original code.¹⁶

Once the MAP point \hat{Z} is reconstructed, we apply the transform to get an unnormalized estimate of the lensing field,

$$\hat{\kappa}_{\log}(\hat{n}) \equiv e^{\hat{Z}(\hat{n})} - \lambda. \quad (44)$$

which we then normalize in the same way as before, by rescaling with the inverse averaged cross-spectrum to input Gaussian fields, Eq. (19). We then reconstruct the biases in the very same way.

The approach we follow (maximizing for Z) does lead to a slightly different map than if we were maximizing directly for κ , even if using the same lognormal prior on Z . This is because the relation between the point-estimates (44) is nonlinear. Our approach is arguably more natural than maximizing for κ , and avoids one issue we encountered in preliminary work, which is how to enforce the lognormality of κ map across iterations: when maximizing for κ , one must ensure that $\kappa(\hat{n}) + \lambda$ is always positive at each point, since $Z(\hat{n})$ is its logarithm.

In Figures 9 we focus on temperature-only reconstructions, and with the LSS non-linear part of the non-Gaussianity only, which has the strongest signature.

For each set of curves, the dashed lines show the biases found in maps with the N-body LSS kappa convergence map, and the dotted lines the biases found in our simple lognormal simulated convergence maps with the same skewness.

In blue and red we have results for the QE and baseline MAP, respectively, where the reconstruction is done

¹⁵ In principle we could also attempt to reconstruct Z jointly with the parameters μ and λ . This is a complication with no relevance for what is being tested here however, and we assume fixed fiducial values of μ and λ , and fixed fiducial spectrum C_L^Z for the reconstructions.

¹⁶ A way to see how the starting point is chosen is the following: if we take the gradient in Z , g_Z , then we require as a first step that: $g_Z \approx g_{Z_0} + F_{Z_0}(Z - Z_0) \approx 0$. By exponentiating the expression that we find for Z we can relate it to $\kappa + \lambda$, and derive that $\hat{Z} = \ln(\hat{\kappa} + \lambda)$.

as in the previous section, with a Gaussian prior on κ . Given the simplicity of the lognormal simulations, it is remarkable how well they reproduce the $N^{(3/2)}$ bias, particularly in the MAP case.

In orange, we show the curves obtained from the reconstructions using the lognormal prior as just discussed. Again, there are essentially no difference found between the N-body input (orange dashed) or lognormal input (orange dotted). Both curves shifts up slightly, reducing slightly the bias on most signal dominated scales, but remain qualitatively very similar.

The fact that the lognormal prior improves only mildly the Gaussian prior is due the fact that we are likelihood dominated, and therefore the MAP reconstruction is mainly data driven.




FIG. 9. Comparison of the $N_L^{(3/2)}$ bias for different variants of temperature reconstructions. The solid lines use the N -body LSS map, while the dashed ones a lognormal map with a similar skewness (at the same input resolution). We show QE results in blue, MAP results with Gaussian prior in red, and MAP results with lognormal prior in orange. We see a small shift upwards between these two choices of priors. The biases found on lognormal maps match well the ones found on the LSS input map. The top shows the bias in the auto-spectrum, and bottom for the cross-spectrum to the true lensing.

B. Bias Hardening

The small-scale CMB temperature is strongly contaminated by foregrounds, limiting the use of more modes to perform CMB lensing reconstruction. Recently, foreground-mitigating lensing reconstruction methods have been developed, allowing more robust and powerful CMB lensing measurements.

In particular, bias-hardening [52, 70, 71], deprojecting the response of the CMB lensing QE to a foregrounds QE, have been used in recent data analyses to extract cosmological parameters from observations.

At the likelihood level, we can derive the bias-hardened QE CMB lensing estimator by looking for modulations in the observed map, going beyond the CMB and noise, that could be attributed to point sources.

We imagine to model our data in pixel space as

$$X^{\text{dat}} = \mathcal{B}\mathcal{D}_\alpha X + \mathcal{B}S + n, \quad (45)$$

where S is some source field. From this equation, the pixel-pixel covariance is

$$C \equiv \langle X^d X^{d,\dagger} \rangle = \text{Cov}_{\bar{\alpha}, S^2} = \mathcal{B}\mathcal{D}_\alpha C^{\text{unl}} \mathcal{D}_\alpha^\dagger \mathcal{B}^\dagger + \mathcal{B}S^2 \mathcal{B}^\dagger + N \quad (46)$$

assuming no cross-terms between CMB, the source, or the noise terms. We can see that at the covariance level the source term induces a variance that is larger than the expected one from just experimental noise.

The log-likelihood becomes then

$$\mathcal{L} \equiv \ln L(X^{\text{dat}} | \bar{\alpha}, S^2) = -\frac{1}{2} X^{\text{dat}} \cdot \text{Cov}_{\bar{\alpha}, S^2}^{-1} X^{\text{dat}} - \frac{1}{2} \det \text{Cov}_{\bar{\alpha}, S^2}. \quad (47)$$

To derive the bias-hardened estimator, we adopt a method akin to the standard quadratic estimator. We start by nulling the gradients for both the lensing potential ϕ and the source term variance S^2 , expanding around $(\phi, S^2) = (0, 0)$ using a first-step Newton iteration:

$$\vec{G}_{\phi, S^2} \approx \vec{G}_{0,0} + F_{0,0} \begin{bmatrix} \phi \\ S^2 \end{bmatrix} \approx 0 \rightarrow \begin{bmatrix} \phi \\ S^2 \end{bmatrix} \approx F_{0,0}^{-1} \vec{G}_{0,0}, \quad (48)$$

where the total gradient of the log-likelihood $\vec{G} = [\frac{\partial \mathcal{L}}{\partial \phi}, \frac{\partial \mathcal{L}}{\partial S^2}]^T$ is approximated by its value at $(\phi, S^2) = (0, 0)$, $\vec{G}_{0,0}$, and a curvature matrix $F_{0,0}$. The latter, encapsulates the responses R^{ab} of the estimator for a to the presence of b , with $a, b \in \phi, S^2$.

In Figures 10, 11 we show results for the $N_L^{(3/2)}$ bias from our simulations using a bias-hardened estimator against a point source. We can see that, compared to the QE, the bias hardening estimator mitigates the impact of the non-Gaussian bias, confirming analytical expectations of [45]. This suggests a bias-hardened MAP-solution might also be helpful, which just adds a prior to Equation 47, which is work in progress.

VI. IMPACT OF $N_L^{(3/2)}$ FOR COSMOLOGICAL ANALYSES

We now turn to the impact of the bias for parameter inference.




FIG. 10. Auto-spectrum $N_L^{(3/2)}$ biases induced by the non-Gaussian LSS-only map on several lensing reconstructions; in dashed we have the QE, in blue, and the bias hardened QE, in purple. The solid line is the MAP. The QE BH is able to mitigate the bias on large scales on a similar level to the MAP estimator, at a small cost in signal to noise compared to the QE. On the left panel, results with temperature data only, while and on the right including polarization too.




FIG. 11. Similar to Figure 10 but for the bias in cross-correlation to the true lensing. After bias-hardening, the quadratic estimator also shows a somewhat reduced $N_L^{(3/2)}$ bias.

A. Bias in the lensing amplitude

In this subsection we estimate the bias in the amplitude of CMB lensing power spectrum due to $N_L^{(3/2)}$. We assume that the covariance matrix of the CMB lensing power spectrum is diagonal and is given by

$$\mathbf{C}_{LL'} = \frac{2\delta_{LL'}}{(2L+1)f_{\text{sky}}} \left(C_L^{\phi\phi} + N_L^{(0)} + N_L^{(1)} \right)^2. \quad (49)$$

We report the signal to noise ratio of the CMB lensing power spectrum between multipoles L_{\min} and L_{\max} as

$$\text{SNR} = \sqrt{\sum_{L_{\min}}^{L_{\max}} C_L^{\phi\phi} \mathbf{C}_{LL'}^{-1} C_{L'}^{\phi\phi}}. \quad (50)$$

This signal to noise ratio is exactly the square root of the Fisher information matrix on the amplitude of the CMB lensing spectrum. The inverse of this SNR thus gives the expected constraint on the amplitude of the CMB lensing spectrum if we only vary this parameter.

We show on the upper panel of the Figure 12 the signal to noise ratio as a function of the maximum lensing scale considered. We see that for the noise levels considered, the polarization estimator brings the most information, but the temperature estimator cannot be neglected. The MAP estimator performs particularly better than the QE especially in the polarization channel. We see that for the polarization, the information gain saturates above the scale $L \sim 1500$, while it still grows up to $L \sim 3000$ for the temperature and minimum variance estimators. This is an important point as future lensing analyses, including cross-correlations with large scale structure, will be able to push the signal to noise to high significances, provided that systematic and modelling treatments are accurately handled.

We estimate the bias in the CMB lensing amplitude as in [72]:

$$b(A_L) = \frac{1}{\text{SNR}^2} \sum_{L_{\min}}^{L_{\max}} \Delta A_L, \quad (51)$$

with

$$\Delta A_L = N_L^{(3/2)} \mathbf{C}_{LL'}^{-1} C_{L'}^{\phi\phi}. \quad (52)$$

We show in the Table II the bias on the lensing amplitude for the different estimators with $L_{\min} = 0$, $L_{\max} = 4000$, including both the LSS and post-Born effects in $N_L^{(3/2)}$. We see that the MAP seems to be almost unbiased on temperature and polarization, but has an almost 1σ bias on the minimum variance combination.

We show in the lower panel of the Figure 12 the bias ΔA_L as a function of scale. As it can be seen from this figure, the $N_L^{(3/2)}$ introduce a scale dependent bias. In some cases, the $N_L^{(3/2)}$ bias flip signs. This lowers the sum in Eq. 51, and thus gives a low bias despite a large absolute value for some given scales. This is the case for the QE in polarization, which has a negative bias for $L \in [300, 1200]$ and positive outside, and for MAP in temperature, which is negative for $L < 1500$ and positive above. This scale dependent bias, which seems to be more important for the MAP than for the QE, could bias the measurements of cosmological parameters that are sensitive to the shape of the lensing spectrum. However, as the total bias almost cancels-out, the parameters combinations sensitive only to the amplitude of the lensing power spectrum should in principle be immune from shape changing effects.

B. Cosmological parameters

We now estimate the biases on cosmological parameters if one does not take into account the $N_L^{(3/2)}$ bias in the analysis. We focus on the sum of the neutrino masses, as this is a key science goal of CMB-S4. We consider the temperature only (TT), the polarization only (Pol), and the combined (MV) estimators, for both QE




FIG. 12. *Upper panel:* Signal to noise ratio on the lensing power spectrum as a function of the maximum scale considered. We show the QE in dashed lines and the MAP in plain lines. The temperature only, polarization only and minimum variance are respectively in blue, orange and green. *Lower panel:* Bias on the lensing power spectrum amplitude as a function of scale, for the different estimators considered, including both the LSS and post-Born effects in $N_L^{(3/2)}$.

Bias	TT	Pol	MV
QE	-3.7σ	0.07σ	-2.13σ
MAP	-0.36σ	-0.12σ	0.93σ

TABLE II. Bias on the lensing power spectrum amplitude, estimated from the Fisher matrix, given in terms of number of sigmas, including both the LSS and post-Born effects in $N_L^{(3/2)}$.

and MAP. Our fiducial cosmology is the Planck FFP10 cosmology¹⁷, with one massive neutrino of 0.06 eV.

Our analysis combines the CMB lensing spectra, the primary (unlensed) CMB spectra and the BAO. We assume that all three sets of observables are independent, so we can sum their log-likelihood. For the BAO we consider a DESI configuration, following the recipes of [73, 74]. For the CMB likelihood, we consider a CMB-

S4 experiment, with a beam $\theta_{\text{FWHM}} = 1$ arcmin, a temperature noise $\Delta T = 1 \mu\text{K-arcmin}$, a polarization noise $\Delta P = \sqrt{2} \mu\text{K-arcmin}$ and a sky fraction $f_{\text{sky}} = 0.4$. For the primary CMB likelihood, we consider unlensed spectra C_ℓ^{TT}, C_ℓ^{TE} and C_ℓ^{EE} , between multipoles 30 and 3000, and assume that these spectra follow a Gaussian likelihood, with Gaussian covariance matrix. We discuss in Appendix B the impact of using the unlensed CMB spectra or the lensed CMB spectra. We show there that the bias on the marginalized cosmological parameters are similar as long as we correctly model the non-Gaussian covariance of the lensed CMB and the correlations between the lensed CMB and the reconstructed lensing potential.

The Gaussian CMB lensing power spectrum likelihood is

$$-2 \ln \mathcal{L}(\theta) = \left(\hat{C}_L^{\phi\phi} - C_L^{\text{th}}(\theta) \right) \mathbf{C}_{LL'}^{-1} \left(\hat{C}_L^{\phi\phi} - C_L^{\text{th}}(\theta) \right), \quad (53)$$

where θ is the set of cosmological parameters being sampled. The covariance matrix is assumed to be diagonal

$$\mathbf{C}_{LL'} = \frac{2\delta_{LL'}}{(2L+1)f_{\text{sky}}} \left(C_L^{\phi\phi} + N_L^{(0)} + N_L^{(1)} \right)^2, \quad (54)$$

where $C_L^{\phi\phi}, N_L^{(0)}$ and $N_L^{(1)}$ are evaluated in the fiducial cosmology. We consider lensing multipoles between 10 and 3000, and assume that the lensing field has been reconstructed with CMB multipoles between 10 and 4000 for both temperature and polarization channels. For the different estimators and different configurations of $N_L^{(3/2)}$ bias (LSS, PB or Total), we generate mock data vectors. These mock data vectors assume that the $N_L^{(0)}$ bias can be perfectly subtracted, and that the $N_L^{(1)}$ cosmology dependence is perfectly modelled. Our mock data vectors are given by

$$\hat{C}_L^{\phi\phi} = C_L^{\phi\phi}(\theta^{\text{fid}}) + N_L^{(1)}(\theta^{\text{fid}}) + \frac{C_L^{\phi\phi}(\theta^{\text{fid}})}{C_L^{\phi\phi, \text{sim}}} N_L^{(3/2)}. \quad (55)$$

where the $N_L^{(3/2)}$ bias has been estimated from simulations, as described in the previous sections, and the CMB lensing power spectrum and $N_L^{(1)}$ bias are taken at the fiducial cosmology. To cancel part of the realization variance of the simulation we divide $N_L^{(3/2)}$ by the lensing power spectrum of the simulation, bin this ratio in 19 multipole bins between 10 and 3810, and fit a spline to these points, weighting by the inverse variance of each bin. We then multiply this spline by our fiducial lensing spectrum. This also allows to partially take into account the difference in the cosmology between the simulation used to estimate the $N_L^{(3/2)}$ (which do not have massive neutrinos, see Equation 13), and the Planck FFP10 cosmology used to generate the data vector. We assume that higher order cosmology dependence of $N_L^{(3/2)}$ can be neglected.

¹⁷ The CAMB parameter file used to generate the spectra can be found in https://github.com/carronj/plancklens/blob/master/plancklens/data/cfs/FFP10_wdipole_params.ini

To accelerate the computation of the theory vector $C_L^{\text{th}}(\theta)$, we do not re-estimate the $N_L^{(1)}$ bias for the sampled cosmology. Instead we correct for the variations of the $N_L^{(1)}$ bias around the fiducial, at first order in $C_L^{\phi\phi}(\theta)$, following the procedure of [19]. The theory vector is then

$$C_L^{\text{th}}(\theta) = C_L^{\phi\phi}(\theta) + N_L^{(1)}(\theta^{\text{fid}}) + \frac{\partial N_L^{(1)}}{\partial C_{L'}^{\phi\phi}} \left[C_{L'}^{\phi\phi}(\theta) - C_{L'}^{\phi\phi}(\theta^{\text{fid}}) \right] \quad (56)$$

where the $N_L^{(1)}$ correction matrix has been previously evaluated in the fiducial cosmology. We neglect the correction due to the variation of the response around the fiducial. We do not expect that this will impact the results. Indeed in the Gaussian case (i.e. without $N_L^{(3/2)}$ bias) we recover unbiased cosmological parameters.

In this analysis we do not model the reconstruction bias $N_L^{(0)}$ in the data vector, we assume it is perfectly subtracted. In a standard analysis one would use the realization dependent bias estimator $\text{RD-}N_L^{(0)}$, which makes the debiasing robust at first order to differences between the fiducial spectra assumed for the reconstruction, and the true CMB spectra of the maps. It was showed in [39, 40] that the realization dependent $\text{RD-}N_L^{(0)}$ allows for unbiased cosmological parameter estimates, for both the QE and the MAP estimators. However we note that this analysis was performed with a Gaussian lensing potential, contrary to the maps we are using here, but we assume that this will have a negligible impact on the $\text{RD-}N_L^{(0)}$ estimate.

We sample for seven cosmological parameters, namely $\ln(10^{10} A_s)$, n_s , θ_{MC} , $\Omega_c h^2$, $\Omega_b h^2$, τ and $\sum m_\nu$. We include a strong Gaussian prior on τ , the reionisation optical depth, assuming a cosmic variance limit of $\sigma_\tau = 0.002$.

We rely on the CAMB cosmological Boltzmann code [75, 76], and sample the posterior with adaptive, speed-hierarchy-aware MCMC sampler (adapted from CosmoMC) [77, 78]. We explore the posterior using the Cobaya [79] and GetDist packages [80].

The marginalized posterior distribution on the sum of the neutrino masses, for our different lensing estimators and source of $N_L^{(3/2)}$ bias, is shown in Figure 13. We show the impact of each term that contributes to the $N_L^{(3/2)}$ bias: the large scale structures non-Gaussianities (LSS) and the post-Born lensing. We confirm that taking into account only one of the two effects, either LSS or PB, will bias the neutrino mass estimates with the QE. We can see that in general, the MAP estimators are less impacted than the QE. When looking at the total $N_L^{(3/2)}$, we see that the MAP is unbiased, both for temperature and for polarization. The QE shows a $\sim 1\sigma$ bias with the temperature reconstruction, while it is unbiased in the polarization.

We note also that as previously observed in the literature [39, 81], the marginalized constraints on the sum of the neutrino mass from the QE is not improved with the MAP. Indeed, even if, as we showed in the Figure 12, the

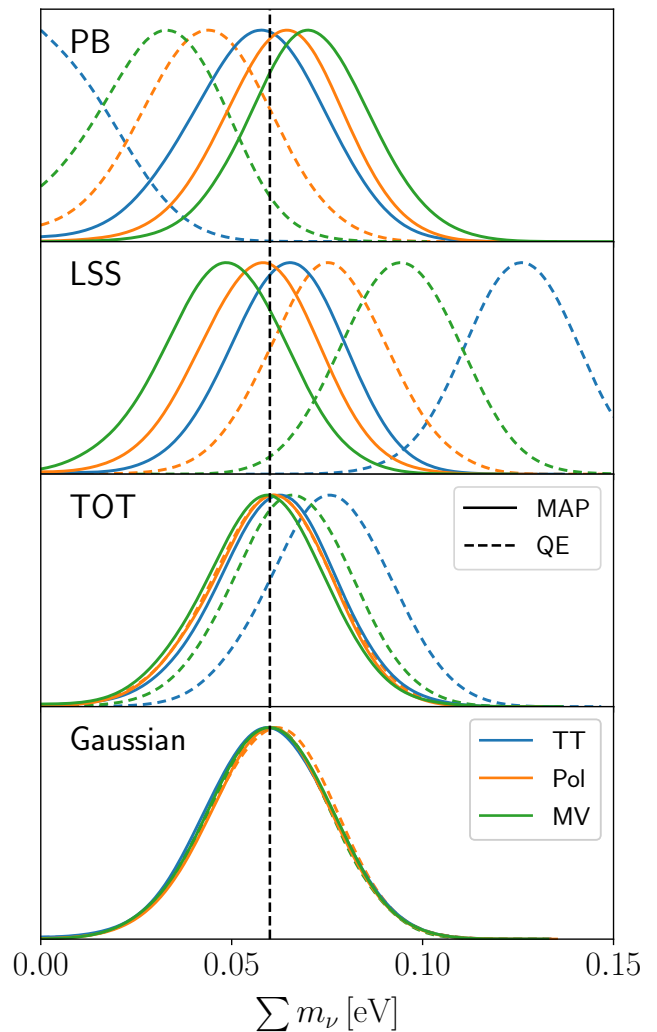


FIG. 13. marginalized posteriors on the sum of the neutrino masses. Panels from top to bottom show respectively the impact of $N_L^{(3/2)}$ when considering the post-Born term only, the large scale structure term only, both terms together, or with no $N_L^{(3/2)}$ bias (fiducial, Gaussian case). The dashed lines are for the QE and the plain lines are for the MAP. Blue, orange and green lines show the temperature only, polarization only or minimum variance estimators. The fiducial cosmology assumes a sum of neutrino masses of 0.06 eV, showed as the vertical dashed line. We combine here the CMB-S4 likelihood with DESI-BAO, and we set a cosmic variance Gaussian prior on the optical depth to reionization.

SNR is increased with the MAP, there are some degeneracies between the cosmological parameters that prevent to reach the full statistical power of the MAP estimator on the marginalized constraints. The marginalized constraints also do not change much if we use the temperature, polarization or the minimum variance estimators, with a one sigma constraint of about 0.016 eV for all cases.

Figure 14 shows the posterior distribution for the n_s , $\Omega_c h^2$ and $\sum m_\nu$ parameters, for the total $N_L^{(3/2)}$ bias




FIG. 14. Posterior distribution for a subset of the sampled cosmological parameters. We show the reconstruction for the minimum variance estimator, for both QE (blue) and MAP (orange), considering the bias due to the total $N_L^{(3/2)}$. The Figure 18 in the Appendix shows the full posterior.

with the minimum variance estimator. We see that even if the MAP is unbiased on the sum of the neutrino mass, it shows a slight 0.5σ bias on the n_s parameter. This bias could be sourced by the shape of the $N_L^{(3/2)}$ bias. Indeed, we showed in Figure 12 that the $N_L^{(3/2)}$ creates a scale dependent bias. In particular, this could tilt the lensing power spectrum, and thus create a bias on the n_s parameter. In comparison, the amplitude bias for the minimum variance QE is mostly negative, so it will not create a strong scale dependent bias, but rather an overall amplitude shift on the CMB lensing power spectrum.

In Appendix B, we compare our results for the QE with the previous results from [42]. They obtained a larger bias on the sum of the neutrino masses. This higher bias comes from the fact that they consider lensed CMB spectra in the likelihood but neglect the non-Gaussian covariance due to lensing and the correlations between CMB and the QE, as described in [82]. In Appendix B we show that using lensed CMB spectra with the correct covariance results in a bias on the sum of the neutrino mass that is similar to the one obtained when considering unlensed spectra and neglecting the non diagonal correlations.

VII. CONCLUSIONS

Future CMB surveys will detect the CMB lensing auto-spectrum well over 100σ significance, and will enable

powerful cosmological constraints on the structure formation history and on the sum of neutrino masses. Great care is required before interpreting these measurements at this high level of sensitivity. Standard quadratic estimators are only guaranteed to be unbiased for purely Gaussian fields and under idealized conditions. When the Gaussianity assumptions fail, future CMB surveys such as CMB-S4 can deliver a biased lensing reconstruction induced by the bispectrum of the lensing potential which, in turn will affect constraints on the sum of neutrino masses [42].

In this work we studied the induced non-Gaussian biases in the auto-spectrum of the CMB lensing potential reconstructed through maximum a posteriori estimators. By using state-of-the art simulations that include non-Gaussian lensing deflections induced by nonlinear matter clustering as well as post-Born lensing corrections, we found that for scales relevant for a CMB-S4 like configuration the non-Gaussian induced bias is mitigated compared to the QE one. When including additional effects due to the post-Born rotation we do not find a significant difference.

When performing a Λ CDM cosmological forecast on the sum of neutrino masses through a full MCMC analysis, we found that contrary to the QE, the MAP estimator does not suffer from biases even if $N_L^{(3/2)}$ biases are not specifically accounted in the fitting of the data vector. This shows that cosmological analyses using a MAP reconstructed lensing potential should in principle be more robust than those based on the standard QE. However, if not correcting for the bias, caution should still be taken when investigating extensions to Λ CDM, as the bias has not fully disappeared, but is only too weak to shift the maximum posterior point, and this could change in other models.

We also tested simple modifications to the MAP estimator to take into account the non-Gaussian statistics of the deflection field. Using the prior of a lognormal field instead of a Gaussian field on the input lensing convergence further mitigates slightly the bias. We have also found that the bias itself found on our N-body simulations is reasonably well reproduced by much simpler lognormal simulations. Generally, differences between reconstructions using the Gaussian or lognormal prior are small, and the lensing reconstructions of similar quality, both on N-body and lognormal simulated inputs. Lognormal simulations are effective in replicating non-Gaussian bias effects from large-scale structure non-linearity. They could be leveraged to predict and subtract biases with sufficient accuracy even when applied to more sophisticated models affecting the full probability density function of the projected matter distribution, such as baryonic effects or extended cosmologies. On the other hand we have found that the lognormal distribution is not a good approximation when including post-Born effects, suggesting that the lognormal approximation is insufficient for recovering the shape of the post-Born induced bispectrum (and same would be

for highly non-linear fields, necessitating alternative approaches [54, 68, 69]). Nevertheless, using state-of-the-art simulations one could learn realistic probability density functions of fields using machine learning techniques to generate new samples or even include alternative priors in our MAP formulation (e.g. [83]).

Our results allow us to conclude that the non-Gaussianity of CMB lensing deflections should not present a major challenge to CMB lensing science in the near future: the bias, small to start with, is mitigated, and the baseline MAP reconstructions with Gaussian prior remain very close to optimal.

In the era of precision-cosmology, additional information will come from the combination and cross-correlation of CMB lensing measurements and multiple large-scale structure observables, such as cosmic shear and galaxy clustering, from DESI, Euclid, and Rubin. It is known that in this case, for QE-based reconstructions, the $N_L^{(3/2)}$ bias in the CMB lensing cross-correlations has a stronger impact depending on the specific tracer and redshift considered in the analysis [45]. This is left for future work.

ACKNOWLEDGMENTS

We thank Antony Lewis for comments. OD, SB, LL and JC acknowledge support from a SNSF Eccellenza Professorial Fellowship (No. 186879). This work was supported by a grant from the Swiss National Supercomputing Centre (CSCS) under project ID s1203. Part of the computations were performed at University of Geneva using Baobab HPC service. GF acknowledges the support of the European Research Council under the Marie Skłodowska Curie actions through the Individual Global Fellowship No. 892401 PiCOGAMBAS.

Appendix A: Mean field impact

The MAP estimator tries to find the best estimate of the CMB lensing potential by nulling the total gradient

$$g^{\text{tot}}(\phi) = g^{\text{QD}}(\phi) + g^{\text{MF}}(\phi) + g^{\text{PR}}(\phi) \quad (\text{A1})$$

Ignoring the mean field contribution g^{MF} , as we often do for simplicity, one is effectively nulling instead

$$g^{\text{tot, noMF}}(\hat{\phi}) = g^{\text{QD}}(\hat{\phi}) + g^{\text{PR}}(\hat{\phi}) \stackrel{!}{=} 0. \quad (\text{A2})$$

The neglect of this gradient piece is equivalent to maximizing a slightly different, but still well-defined likelihood function. Hence the iterative procedure converges without problems. The actual total gradient at this point $\hat{\phi}$ is not zero but

$$g^{\text{tot}}(\hat{\phi}) = g^{\text{MF}}(\hat{\phi}) \neq 0. \quad (\text{A3})$$

Early indications from reconstructions from polarization suggested this term was small [47]. Physically, this is because this term accounts for the anisotropies in the noise maps induced by delensing, and these anisotropies trace for the most part the magnification part of the lensing signal, rather than the shear-like signal that the most powerful EB polarized estimator is sensitive to. Delensing affects the local noise levels by changing areas according to the local magnification, given by $1 - 2\kappa$ to first order, and the mean-field removes the contribution of this noise anisotropy to the quadratic gradient piece.

It turns out that that for reconstructions from temperature, that takes substantial contribution from the magnification-like lensing signals, the mean field contribution is much larger. If the mean-field at the solution $\hat{\kappa}_{LM}$ were exactly proportional to $\hat{\kappa}_{LM}$ as suggested by the argument above, the mean-field contamination of $\hat{\kappa}$ acts simply as a rescaling of the output map. This would have no impact on the cross-correlation coefficient of the reconstruction with the input (the ‘quality’ of the reconstruction), only on its normalization

$$W_L^{\text{emp}} \equiv \frac{C_L^{\hat{\phi}\phi}}{C_L^{\phi\phi}} \quad (\text{A4})$$

(the empirical ‘Wiener-filter’ of the reconstruction). The shift seen on temperature reconstruction for our CMB-S4-like configuration is shown on Fig. 15, and is substantial. Not accounting for the mean field lowers the Wiener-filter curve. This is because the magnification follows κ but with a minus sign. In Fig. 16 we show that




FIG. 15. Comparison of the empirical normalization of output lensing maps, after a few iterations, for our temperature-based CMB-S4 like reconstructions. In blue we show the case of the QE estimator, in green for the MAP estimator if ignoring the mean field contribution, and in orange when including it, using an estimate from a finite number of Monte-Carlo at each step.

after applying the empirical normalization, for the case of mean-field subtraction/no subtraction, we find a power

spectrum in line with naive expectations,

$$C_L^{\hat{\phi}\hat{\phi}} \sim C_L^{\phi\phi} + N_L^{(0)} + N_L^{(1)} \quad (\text{A5})$$

where the lensing biases are calculated using the partially delensed CMB spectra. The change in cross-correlation coefficient between the two maps is tiny. This gives support to the idea that the bulk of the effect is only a rescaling. Hence, in this paper we proceed using the empirical normalization, and neglecting the delensed-noise mean-field altogether. A more careful study of the mean-field is ongoing.




FIG. 16. Lensing spectrum from a MAP temperature reconstruction, after dividing by the empirical normalization of the lensing map (blue). In red the naive theory prediction, obtained from a iterative $N_L^{(0)}$ and $N_L^{(1)}$ calculation. We get a similar result if we account for the mean field, justifying the usage of an empirical normalisation when not account for a mean-field.

Appendix B: About the impact of non-Gaussian lensing correlations on the cosmological parameter biases

In the previous study of [42] (hereafter Beck18), the likelihood was neglecting the correlations between the reconstructed lensing field and the CMB spectra. In the results we present in the Section VIB above, assume we can neglect these correlations by using unlensed CMB spectra un the likelihood. We will now make systematic study of this assumption, and evaluate the impact of using the lensed or unlensed CMB spectra, and the impact of the non-Gaussian covariance on the final cosmological constraints.

We will consider four different likelihood configurations:

1. The Beck18 likelihood
2. The Beck18 likelihood but with unlensed CMB spectra

3. The likelihood we introduced in Section VIB, which uses unlensed CMB spectra and assumes no correlations between the lensing power spectrum and the CMB fields

4. A complete likelihood including the full non-Gaussian correlations between the lensed CMB spectra and the lensing potential, following [82]

The Beck18 CMB likelihood is

$$-2 \log \mathcal{L}(\theta|\hat{\mathbf{C}}) = \sum_{\ell} (2\ell + 1) f_{\text{sky}} \left(\ln \frac{|\mathbf{C}_{\ell}|}{|\hat{\mathbf{C}}_{\ell}|} + \mathbf{C}_{\ell}^{-1} \hat{\mathbf{C}}_{\ell} - 3 \right) \quad (\text{B1})$$

where the theoretical covariance matrix \mathbf{C} and the data covariance matrix $\hat{\mathbf{C}}$ are constructed with the lensed CMB spectra and are given by

$$\mathbf{C}_{\ell} = \begin{pmatrix} C_{\ell}^{TT} + N_{\ell}^{TT} & C_{\ell}^{TE} & C_{\ell}^{T\phi} \\ C_{\ell}^{TE} & C_{\ell}^{EE} + N_{\ell}^{EE} & 0 \\ C_{\ell}^{T\phi} & 0 & C_{\ell}^{\phi\phi} + N_{\ell}^{\phi\phi} \end{pmatrix} \quad (\text{B2})$$

For the theoretical covariance \mathbf{C}_{ℓ} , the CMB lensing noise is $N_{\ell}^{\phi\phi} = N_{\ell}^{(0)} + N_{\ell}^{(1)}$. For the mock data covariance we include the $N_L^{(3/2)}$ bias estimated from the simulations as

$$N_L^{\phi\phi} = N_L^{(0)} + N_L^{(0)} + \frac{C_L^{\phi\phi, \text{fid}}}{C_L^{\phi\phi, \text{sim}}} N_L^{(3/2), \text{sim}}. \quad (\text{B3})$$

In this configuration, and in order to reproduce the results of Beck18, we do not take into account the cosmological dependence of the $N_L^{(0)}$ and $N_L^{(1)}$ biases, and we keep them fixed to their fiducial value.

The Beck18 likelihood with unlensed spectra is the same as above but replaces the lensed CMB spectra by their unlensed version.

The likelihood which takes into account the non-Gaussian covariance introduced in [82] has the following data vector:

$$\mathcal{C}_{\ell} = \left(C_{\ell}^{TT}, C_{\ell}^{EE}, C_{\ell}^{TE}, C_{\ell}^{\phi\phi} + N_{\ell}^{(1)} \right). \quad (\text{B4})$$

For the theory vector $\mathcal{C}_{\ell}(\theta)$ we vary the theoretical $N_{\ell}^{(1)}$ at first order in the CMB lensing spectra when sampling the cosmology, like in Eq. 56. The mock data vector $\hat{\mathcal{C}}_{\ell}$ includes the $N_L^{(3/2)}$ bias, which is not modelled by the theory vector. The full likelihood is then

$$-2 \log \mathcal{L}(\theta) = \left(\hat{\mathcal{C}}_{\ell} - \mathcal{C}_{\ell}(\theta) \right) \mathbf{C}_{\ell\ell'}^{-1} \left(\hat{\mathcal{C}}_{\ell'} - \mathcal{C}_{\ell'}(\theta) \right) \quad (\text{B5})$$

The covariance matrix $\mathbf{C}_{\ell\ell'}$ includes the non-Gaussian terms described in [82], such as the correlations between the CMB spectra due to lensing, and the correlation between the reconstructed lensing field and the CMB.

In all likelihood scenarios we generate the data vector using the same cosmology of Beck18, reproduced in Table B, contrary to the analysis in the Section VIB where

$\Omega_b h^2$	0.02225 ± 0.00016
$\Omega_c h^2$	0.1198 ± 0.0015
τ	0.058 ± 0.012
$\log(10^{10} A_s)$	3.094 ± 0.034
n_s	0.9645 ± 0.0049
$100\theta_{MC}$	1.04077 ± 0.00032
$\sum m_\nu$ [eV]	[0, 300]

TABLE III. Fiducial cosmological parameters used in this section, following Beck18, together with their 1σ Gaussian prior, or uniform parameter bound.

we used the FFP10 cosmology. Notably, we now consider massless neutrinos, instead of the minimal mass normal hierarchy considered before. Moreover, contrary to the main analysis, the likelihoods used here do not include external BAO constraints. We follow again Beck18 and assume tight Gaussian priors on the parameters, except for the sum of neutrino mass which assumes a flat prior, given on Table B.

We perform the MCMC samplings like in the main analysis, and we now discuss the results obtained for the QE minimum variance estimator, with the LSS only term in the $N_L^{(3/2)}$ bias. We show in the Figure 17 the constraints for three cosmological parameters $\Omega_c h^2$, $\ln(10^{10} A_s)$ and $\sum m_\nu$. The purple contours are the one obtained with the likelihood of Beck18, with the lensed CMB spectra. We retrieve the same results as in Beck18 (see their Figure 11), with a posterior estimate on the sum of neutrino masses peaking at 0.18 eV. The red contours show the posterior with the Beck18 likelihood with unlensed CMB spectra in the data and theory vectors. The blue contours are for our likelihood, using unlensed CMB spectra, for Beck18 cosmology and priors. In practice, the main difference here with the Beck18 unlensed likelihood is that we take into account the variation of the $N_L^{(1)}$ bias when sampling the cosmology, while it is not the case for the Beck18 likelihood. Finally, the green contours are for the likelihood with lensed CMB spectra and the [82] non-Gaussian covariance.

It appears that the Beck18 likelihood, with the lensed CMB spectra, has the largest bias on the sum of the neutrino masses and on the amplitude of the matter power spectrum. We see that when considering unlensed CMB

spectra the Beck18 likelihood obtains similar constraints as ours, which has unlensed spectra as well. We show that we obtain almost exactly the same posterior, with a reduced bias, if we use the lensed spectra but correctly taking into account the correlations due to lensing from [82]. This reduction of the bias when using the full non-Gaussian covariance can be understood as lowering the impact of the bias from the lensing power spectrum, by avoiding to “double” count the biased lensing power spectrum in the analysis. If we include as well the lensed C_ℓ^{BB} spectrum in the data vector, even if the correlations are properly taken into account, the bias is not reduced. We interpret that as the fact that the lensed spectrum is another measurement of $C_L^{\phi\phi}$, as we can approximate for $\ell \ll 1000$ [3]

$$C_\ell^{BB} \sim \frac{1}{4\pi} \int \frac{d\ell'}{\ell'} \ell'^4 C_{\ell'}^{\phi\phi} \ell'^2 C_{\ell'}^{EE, \text{unl}}. \quad (\text{B6})$$

Thus there is a tension between the lensed BB spectrum, which measures an unbiased $C_L^{\phi\phi}$, and the estimate lensing power spectrum which is biased by $N_L^{(3/2)}$. It appears that the biased $C_L^{\phi\phi}$ dominates and the likelihood is leaning towards the biased measurement of the neutrino mass. In practice, future surveys might not include the lensed BB spectrum in the analysis since the BB spectrum might be more prone to instrumental systematics and polarized foregrounds such as point sources. But comparing the marginalized constraints with or without the BB spectra in the likelihood could serve as additional robustness test to assess the presence of biases coming from non-Gaussian effects.

We do not test the impact of the full non-Gaussian covariance for the MAP estimator. In that case, one would use the partially delensed CMB spectra in the likelihood, reducing the non-diagonal terms in the covariance matrix. Fisher forecasts using the covariance matrix computed with delensed spectra were published in [93, 94]. We leave for a future work a detailed comparison of our MAP reconstructed delensed spectra with this delensed analytical covariance.

We show in the Figure 18 the full posterior of the likelihood for the QE and MAP reconstruction, considering the total $N_L^{(3/2)}$ bias, and seven cosmological parameters as described in the Section VI B.

* omar.darwish@unige.ch

† sebastian.belkner@unige.ch

- [1] J. Aguirre *et al.* (Simons Observatory), The Simons Observatory: Science goals and forecasts, *JCAP* **1902**, 056 (2019), arXiv:1808.07445 [astro-ph.CO].
- [2] K. N. Abazajian *et al.*, CMB-S4 Science Book, First Edition, (2016), arXiv:1610.02743.
- [3] A. Lewis and A. Challinor, Weak gravitational lensing of the CMB, *Physics Reports* **429**, 1 (2006), arXiv:astro-ph/0601594.
- [4] R. de Putter, O. Zahn, and E. V. Linder, CMB lensing constraints on neutrinos and dark energy, *Phys. Rev. D* **79**, 065033 (2009), arXiv:0901.0916.
- [5] R. Stompor and G. Efstathiou, Gravitational lensing of cosmic microwave background anisotropies and cosmological parameter estimation, *MNRAS* **302**, 735 (1999), arXiv:astro-ph/9805294 [astro-ph].

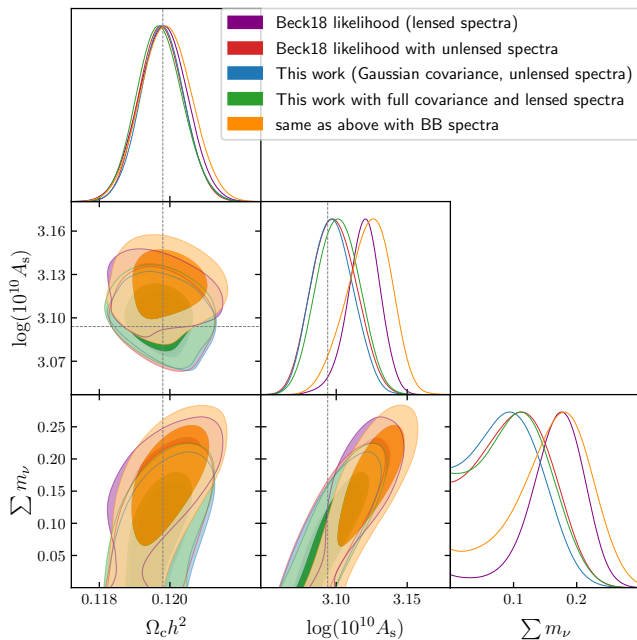


FIG. 17. marginalized contours for three cosmological parameters, for the QE minimum variance, in the presence of the LSS only $N_L^{(3/2)}$ bias. We show the results for the Beck18 likelihood with the lensed CMB spectra (in purple) or with the unlensed CMB spectra (in red). We compare with our likelihood used for the main results in Section VIB (in blue). We also show in green the likelihood which includes the full non-Gaussian covariance of the lensed CMB spectra, and the correlations between the reconstructed lensing spectrum and the CMB. We see that correctly taking into account these correlations gives almost the same posterior as considering unlensed CMB spectra and neglecting correlations. In orange we add the lensed BB spectrum in the data vector, taking into account its correlations with the lensed CMB and the reconstructed lensing potential.

- [6] N. Sailer, E. Castorina, S. Ferraro, and M. White, Cosmology at high redshift — a probe of fundamental physics, *Journal of Cosmology and Astroparticle Physics* **2021** (12), 049.
- [7] A. Vallinotto, Using Cosmic Microwave Background Lensing to Constrain the Multiplicative Bias of Cosmic Shear, *Astrophys. J.* **759**, 32 (2012), arXiv:1110.5339.
- [8] E. Schaan, E. Krause, T. Eifler, O. Doré, H. Miyatake, J. Rhodes, and D. N. Spergel, Looking through the same lens: Shear calibration for LSST, Euclid, and WFIRST with stage 4 CMB lensing, *Phys. Rev. D* **95**, 123512 (2017), arXiv:1607.01761.
- [9] R. Cawthon, Effects of redshift uncertainty on cross-correlations of CMB lensing and galaxy surveys, *Phys. Rev. D* **101**, 063509 (2020), arXiv:1809.09251.
- [10] A. Vallinotto, The Synergy between the Dark Energy Survey and the South Pole Telescope, *Astrophys. J.* **778**, 108 (2013), arXiv:1304.3474.
- [11] L. Wenzl, C. Doux, C. Heinrich, R. Bean, B. Jain, O. Doré, T. Eifler, and X. Fang, Cosmology with the Roman Space Telescope - Synergies with CMB lensing, *MNRAS* **512**, 5311 (2022), arXiv:2112.07681.
- [12] Euclid Collaboration, S. Ilić, *et al.*, Euclid preparation. XV. Forecasting cosmological constraints for the Euclid and CMB joint analysis, *A&A* **657**, A91 (2022), arXiv:2106.08346.
- [13] A. Lewis and L. King, Cluster masses from CMB and galaxy weak lensing, *Phys. Rev. D* **73**, 063006 (2006), arXiv:astro-ph/0512104 [astro-ph].
- [14] W. Hu, S. DeDeo, and C. Vale, Cluster mass estimators from CMB temperature and polarization lensing, *New Journal of Physics* **9**, 441 (2007), arXiv:astro-ph/0701276 [astro-ph].
- [15] S. Raghunathan, S. Patil, E. J. Baxter, F. Bianchini, L. E. Bleem, T. M. Crawford, G. P. Holder, A. Manzotti, and C. L. Reichardt, Measuring galaxy cluster masses with CMB lensing using a Maximum Likelihood estimator: statistical and systematic error budgets for future experiments, *JCAP* **2017**, 030 (2017), arXiv:1705.00411.
- [16] U. Seljak and C. M. Hirata, Gravitational lensing as a contaminant of the gravity wave signal in the CMB, *Phys. Rev. D*




FIG. 18. Full posterior distribution for the seven sampled cosmological parameters. We show the reconstruction for the minimum variance estimator, for both QE (blue) and MAP (orange), considering the bias due to the total $N_L^{(3/2)}$.

- 69**, 043005 (2004), arXiv:astro-ph/0310163 [astro-ph].
- [17] K. M. Smith, D. Hanson, M. LoVerde, C. M. Hirata, and O. Zahn, Delensing CMB polarization with external datasets, *JCAP* **2012**, 014 (2012), arXiv:1010.0048.
- [18] W. Hu and T. Okamoto, Mass Reconstruction with Cosmic Microwave Background Polarization, *Astrophys. J.* **574**, 566 (2002), arXiv:astro-ph/0111606.
- [19] N. Aghanim *et al.* (Planck), Planck 2018 results. VIII. Gravitational lensing, *A&A* **641**, A8 (2020), arXiv:1807.06210 [astro-ph.CO].
- [20] J. Carron, M. Mirmelstein, and A. Lewis, CMB lensing from Planck PR4 maps, *JCAP* **2022**, 039 (2022), arXiv:2206.07773.
- [21] M. S. Madhavacheril *et al.*, The Atacama Cosmology Telescope: DR6 Gravitational Lensing Map and Cosmological Parameters, (2023), arXiv:2304.05203.
- [22] W. L. K. Wu *et al.*, A Measurement of the Cosmic Microwave Background Lensing Potential and Power Spectrum from 500 deg² of SPTpol Temperature and Polarization Data, *Astrophys. J.* **884**, 70 (2019), arXiv:1905.05777.
- [23] M. A. Faúndez *et al.*, Measurement of the Cosmic Microwave Background Polarization Lensing Power Spectrum from Two Years of POLARBEAR Data, *Astrophys. J.* **893**, 85 (2020), arXiv:1911.10980.
- [24] BICEP/Keck Collaboration, :, P. A. R. Ade, *et al.*, BICEP / Keck XVII: Line of Sight Distortion Analysis: Estimates of Gravitational Lensing, Anisotropic Cosmic Birefringence, Patchy Reionization, and Systematic Errors, (2022), arXiv:2210.08038.
- [25] N. Sailer, S. Ferraro, and E. Schaan, Foreground-immune CMB lensing reconstruction with polarization, *Phys. Rev. Lett.* **107**, 023504 (2023), arXiv:2211.03786.
- [26] A. S. Maniyar, Y. Ali-Haïmoud, J. Carron, A. Lewis, and M. S. Madhavacheril, Quadratic estimators for CMB weak

- lensing, *Phys. Rev. D* **103**, 083524 (2021), arXiv:2101.12193.
- [27] J. Carron, Real-world CMB lensing quadratic estimator power spectrum response, *JCAP* **02**, 057 (2023).
- [28] D. Beck, J. Errard, and R. Stompor, Impact of Polarized Galactic Foreground Emission on CMB Lensing Reconstruction and Delensing of B-Modes, *JCAP* **06**, 030 (2020), arXiv:2001.02641.
- [29] E. Schaan and S. Ferraro, Foreground-Immune Cosmic Microwave Background Lensing with Shear-Only Reconstruction, *Phys. Rev. Lett.* **122**, 181301 (2019), arXiv:1804.06403.
- [30] M. Mirmelstein, G. Fabbian, A. Lewis, and J. Peloton, Instrumental systematics biases in CMB lensing reconstruction: A simulation-based assessment, *Phys. Rev. D* **103**, 123540 (2021), arXiv:2011.13910.
- [31] O. Darwish, B. D. Sherwin, N. Sailer, E. Schaan, and S. Ferraro, Optimizing foreground mitigation for cmb lensing with combined multifrequency and geometric methods, *Phys. Rev. D* **107**, 043519 (2023), arXiv:2111.00462.
- [32] M. Millea, E. Anderes, and B. D. Wandelt, Bayesian delensing of CMB temperature and polarization, *Phys. Rev. D* **100**, 023509 (2019).
- [33] M. Millea, E. Anderes, and B. D. Wandelt, Sampling-based inference of the primordial CMB and gravitational lensing, *Phys. Rev. D* **102**, 123542 (2020).
- [34] M. Millea and U. Seljak, Marginal unbiased score expansion and application to CMB lensing, *Phys. Rev. D* **105**, 103531 (2022).
- [35] C. M. Hirata and U. Seljak, Reconstruction of lensing from the cosmic microwave background polarization, *Phys. Rev. D* **68**, 083002 (2003).
- [36] J. Carron and A. Lewis, Maximum a posteriori CMB lensing reconstruction, *Phys. Rev. D* **96**, 063510 (2017).
- [37] S. Adachi, others, and Polarbear Collaboration, Internal Delensing of Cosmic Microwave Background Polarization B-Modes with the POLARBEAR Experiment, *Phys. Rev. Lett.* **124**, 131301 (2020), arXiv:1909.13832.
- [38] M. Millea *et al.*, Optimal Cosmic Microwave Background Lensing Reconstruction and Parameter Estimation with SPTpol Data, *Astrophys. J.* **922**, 259 (2021), arXiv:2012.01709.
- [39] L. Legrand and J. Carron, Lensing power spectrum of the cosmic microwave background with deep polarization experiments, *Phys. Rev. D* **105**, 123519 (2022), arXiv:2112.05764.
- [40] L. Legrand and J. Carron, Robust and efficient CMB lensing power spectrum from polarization surveys, *Phys. Rev. D* **108**, 103516 (2023), arXiv:2304.02584 [astro-ph.CO].
- [41] V. Böhm, M. Schmittfull, and B. D. Sherwin, A bias to CMB lensing measurements from the bispectrum of large-scale structure, **94**, 043519 (), 1605.01392.
- [42] D. Beck, G. Fabbian, and J. Errard, Lensing Reconstruction in Post-Born Cosmic Microwave Background Weak Lensing, **98**, 043512, 1806.01216.
- [43] V. Böhm, B. D. Sherwin, J. Liu, J. C. Hill, M. Schmittfull, and T. Namikawa, On the effect of non-Gaussian lensing deflections on CMB lensing measurements, **98**, 123510 (), 1806.01157.
- [44] G. Pratten and A. Lewis, Impact of post-Born lensing on the CMB, *JCAP* **8**, 047 (2016), arXiv:1605.05662.
- [45] G. Fabbian, A. Lewis, and D. Beck, CMB lensing reconstruction biases in cross-correlation with large-scale structure probes, **2019**, 057, 1906.08760.
- [46] J. Carron and A. Lewis, Maximum a posteriori CMB lensing reconstruction, **96**, 063510, 1704.08230.
- [47] S. Belkner, J. Carron, L. Legrand, C. Umiltà, C. Pryke, and C. Bischoff (CMB-S4), CMB-S4: Iterative internal delensing and r constraints, (2023), arXiv:2310.06729 [astro-ph.CO].
- [48] C. M. Hirata and U. Seljak, Analyzing weak lensing of the cosmic microwave background using the likelihood function, **67**, 043001 ().
- [49] C. M. Hirata and U. Seljak, Reconstruction of lensing from the cosmic microwave background polarization, **68**, 083002 ().
- [50] D. Hanson, A. Challinor, and A. Lewis, Weak lensing of the CMB, *General Relativity and Gravitation* **42**, 2197 (2010).
- [51] M. Kesden, A. Cooray, and M. Kamionkowski, Lensing reconstruction with cmb temperature and polarization, *Physical Review D* **67**, 10.1103/physrevd.67.123507 (2003).
- [52] T. Namikawa, D. Hanson, and R. Takahashi, Bias-hardened cmb lensing, *Monthly Notices of the Royal Astronomical Society* **431**, 609–620 (2013).
- [53] M. Reinecke, S. Belkner, and J. Carron, Improved cosmic microwave background (de-)lensing using general spherical harmonic transforms (2023), arXiv:2304.10431 [astro-ph.CO].
- [54] G. Fabbian, M. Calabrese, and C. Carbone, CMB weak-lensing beyond the Born approximation: a numerical approach, *JCAP* **2**, 050 (2018), arXiv:1702.03317.
- [55] S. Hilbert, J. Hartlap, S. D. M. White, and P. Schneider, Ray-tracing through the Millennium Simulation: Born corrections and lens-lens coupling in cosmic shear and galaxy-galaxy lensing, *A&A* **499**, 31 (2009), arXiv:0809.5035.
- [56] C. Carbone, M. Petkova, and K. Dolag, DEMNUni: ISW, Rees-Sciama, and weak-lensing in the presence of massive neutrinos, *JCAP* **7**, 034 (2016), arXiv:1605.02024.
- [57] E. Castorina, C. Carbone, J. Bel, E. Sefusatti, and K. Dolag, DEMNUni: the clustering of large-scale structures in the presence of massive neutrinos, *JCAP* **7**, 043 (2015), arXiv:1505.07148.
- [58] C. Carbone, V. Springel, C. Baccigalupi, M. Bartelmann, and S. Matarrese, Full-sky maps for gravitational lensing of the cosmic microwave background, *MNRAS* **388**, 1618 (2008), arXiv:0711.2655.
- [59] M. Calabrese, C. Carbone, G. Fabbian, M. Baldi, and C. Baccigalupi, Multiple lensing of the cosmic microwave background anisotropies, *JCAP* **3**, 049 (2015), arXiv:1409.7680.
- [60] P. Fosalba, E. Gaztañaga, F. J. Castander, and M. Manera, The onion universe: all sky lightcone simulations in spherical shells, *MNRAS* **391**, 435 (2008), arXiv:0711.1540.
- [61] G. Fabbian and R. Stompor, High-precision simulations of the weak lensing effect on cosmic microwave background polar-

- ization, *A&A* **556**, A109 (2013), arXiv:1303.6550 [astro-ph.CO].
- [62] H. S. Xavier, F. B. Abdalla, and B. Joachimi, Improving lognormal models for cosmological fields, **459**, 3693, 1602.08503.
- [63] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann, HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere, *Astrophys. J.* **622**, 759 (2005), arXiv:astro-ph/0409513 [astro-ph].
- [64] A. Zonca, L. P. Singer, D. Lenz, M. Reinecke, C. Rosset, E. Hivon, and K. M. Gorski, healpy: equal area pixelization and spherical harmonics transforms for data on the sphere in python, *Journal of Open Source Software* **4**, 1298 (2019).
- [65] A. Hall and A. Taylor, Non-Gaussian likelihood of weak lensing power spectra, *Phys. Rev. D* **105**, 123527 (2022), arXiv:2202.04095 [astro-ph.CO].
- [66] H. Gil-Marín, C. Wagner, F. Fragkoudi, R. Jimenez, and L. Verde, An improved fitting formula for the dark matter bispectrum, **2012**, 047, 1111.4477.
- [67] T. Namikawa, B. Bose, F. R. Bouchet, R. Takahashi, and A. Taruya, CMB lensing bi-spectrum: Assessing analytical predictions against full-sky lensing simulations, **99**, 063511, 1812.10635.
- [68] S. Das and J. P. Ostriker, Testing a New Analytic Model for Gravitational Lensing Probabilities, *Astrophys. J.* **645**, 1 (2006), arXiv:astro-ph/0512644 [astro-ph].
- [69] A. Barthelemy, S. Codis, and F. Bernardeau, Post-Born corrections to the one-point statistics of (CMB) lensing convergence obtained via large deviation theory, *MNRAS* **494**, 3368 (2020), arXiv:2002.03625 [astro-ph.CO].
- [70] S. J. Osborne, D. Hanson, and O. Doré, Extragalactic foreground contamination in temperature-based cmb lens reconstruction, *Journal of Cosmology and Astroparticle Physics* **2014** (03), 024–024.
- [71] N. Sailer, E. Schaan, and S. Ferraro, Lower bias, lower noise cmb lensing with foreground-hardened estimators, *Physical Review D* **102**, 10.1103/physrevd.102.063517 (2020).
- [72] A. Amara and A. Refregier, Systematic Bias in Cosmic Shear: Beyond the Fisher Matrix, *Mon. Not. Roy. Astron. Soc.* **391**, 228 (2008), arXiv:0710.5171 [astro-ph].
- [73] A. Font-Ribera, P. McDonald, N. Mostek, B. A. Reid, H.-J. Seo, and A. Slosar, DESI and other dark energy experiments in the era of neutrino mass measurements, *JCAP* **05**, 023, arXiv:1308.4164 [astro-ph.CO].
- [74] A. Aghamousa *et al.* (DESI), The DESI Experiment Part I: Science, Targeting, and Survey Design, (2016), arXiv:1611.00036 [astro-ph.IM].
- [75] A. Lewis, A. Challinor, and A. Lasenby, Efficient computation of CMB anisotropies in closed FRW models, *Astrophys. J.* **538**, 473 (2000), arXiv:astro-ph/9911177 [astro-ph].
- [76] C. Howlett, A. Lewis, A. Hall, and A. Challinor, CMB power spectrum parameter degeneracies in the era of precision cosmology, *JCAP* **1204**, 027, arXiv:1201.3654 [astro-ph.CO].
- [77] A. Lewis and S. Bridle, Cosmological parameters from CMB and other data: A Monte Carlo approach, *Phys. Rev.* **D66**, 103511 (2002), arXiv:astro-ph/0205436 [astro-ph].
- [78] A. Lewis, Efficient sampling of fast and slow cosmological parameters, *Phys. Rev.* **D87**, 103529 (2013), arXiv:1304.4473 [astro-ph.CO].
- [79] J. Torrado and A. Lewis, Cobaya: Code for Bayesian Analysis of hierarchical physical models, *JCAP* **05**, 057, arXiv:2005.05290 [astro-ph.IM].
- [80] A. Lewis, GetDist: a Python package for analysing Monte Carlo samples, (2019), arXiv:1910.13970 [astro-ph.IM].
- [81] R. Allison, P. Caulal, E. Calabrese, J. Dunkley, and T. Louis, Towards a cosmological neutrino mass detection, *Phys. Rev. D* **92**, 123535 (2015), arXiv:1509.07471 [astro-ph.CO].
- [82] J. Peloton, M. Schmittfull, A. Lewis, J. Carron, and O. Zahn, Full covariance of CMB and lensing reconstruction power spectra, *Phys. Rev. D* **95**, 043508 (2017), arXiv:1611.01446 [astro-ph.CO].
- [83] A. Rouhiainen, U. Giri, and M. Münchmeyer, Normalizing flows for random fields in cosmology (2021), arXiv:2105.12024.
- [84] R. M. Neal, Taking Bigger Metropolis Steps by Dragging Fast Variables, *ArXiv Mathematics e-prints* (2005), math/0502099.
- [85] V. Böhm, C. Modi, and E. Castorina, Lensing corrections on galaxy-lensing cross correlations and galaxy-galaxy auto correlations, **2020**, 045 (), 1910.06722.
- [86] S. S. Boruah, E. Roza, and P. Fiedorowicz, Map-based cosmology inference with lognormal cosmic shear maps, 2204.13216.
- [87] A. Buchalter, M. Kamionkowski, and A. H. Jaffe, The Angular Three-Point Correlation Function in the Quasilinear Regime, **530**, 36, astro-ph/9903486.
- [88] P. Coles and B. Jones, A lognormal model for the cosmological mass distribution, **248**, 1.
- [89] S. Hilbert, J. Hartlap, and P. Schneider, Cosmic shear covariance: The log-normal approximation, **536**, A85.
- [90] A. Lewis and G. Pratten, Effect of lensing non-Gaussianity on the CMB power spectra, *JCAP* **12**, 003 (2016), arXiv:1608.01263.
- [91] C. Carbone, C. Baccigalupi, M. Bartelmann, S. Matarrese, and V. Springel, Lensed CMB temperature and polarization maps from the Millennium Simulation, *MNRAS* **396**, 668 (2009), arXiv:0810.4145.
- [92] M. R. Becker, CALCLENS: weak lensing simulations for large-area sky surveys and second-order effects in cosmic shear power spectra, *MNRAS* **435**, 115 (2013), arXiv:arXiv:1210.3069.
- [93] S. C. Hotinli, J. Meyers, C. Trendafilova, D. Green, and A. van Engelen, The benefits of CMB delensing, *JCAP* **04** (04), 020, arXiv:2111.15036 [astro-ph.CO].
- [94] D. Green, J. Meyers, and A. van Engelen, CMB Delensing Beyond the B Modes, *JCAP* **12**, 005, arXiv:1609.08143 [astro-ph.CO].
- [95] Y. Fantaye, C. Baccigalupi, S. Leach, and A. P. S. Yadav, CMB lensing reconstruction in the presence of diffuse polarized foregrounds, *JCAP* **12**, 017 (2012), arXiv:1207.0508.

- [96] E. Schaan, S. Ferraro, and U. Seljak, Photo-z outlier self-calibration in weak lensing surveys, *JCAP* **2020**, 001 (2020), arXiv:2007.12795.