

Lensing power spectrum of the cosmic microwave background with deep polarization experiments

Louis Legrand^{✉*} and Julien Carron^{✉†}

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

 (Received 14 December 2021; accepted 3 June 2022; published 17 June 2022)

Precise reconstruction of the cosmic microwave background lensing potential can be achieved with deep polarization surveys by iteratively removing lensing-induced B modes. We introduce a lensing spectrum estimator and its likelihood for such optimal iterative reconstruction. Our modeling shares similarities to the state-of-the-art likelihoods for quadratic estimator (QE)-based lensing reconstruction. In particular, we generalize the $N^{(0)}$ and $N^{(1)}$ lensing biases, and design a realization-dependent spectrum debiaser, making this estimator robust to uncertainties in the data modeling. We demonstrate unbiased recovery of the cosmology using map-based reconstructions, focussing on lensing-only cosmological constraints and neutrino mass measurement in combination with CMB spectra and acoustic oscillation data. We find this spectrum estimator is essentially optimal and with a diagonal covariance matrix. For a CMB-S4 survey, this likelihood can double the constraints on the lensing amplitude compared to the QE on a wide range of scales, while at the same time keeping numerical cost under control and being robust to errors.

DOI: [10.1103/PhysRevD.105.123519](https://doi.org/10.1103/PhysRevD.105.123519)

I. INTRODUCTION

By probing the matter density fluctuations up to the last scattering surface, gravitational lensing of the cosmic microwave background (CMB) is a powerful probe of the growth of structures, able to constrain the Λ CDM cosmological model, the neutrino mass scale, or modified gravity theories [1]. State-of-the-art reconstructions of CMB lensing maps [2–5] have been obtained with a quadratic estimator (QE) [6], which uses the anisotropic signatures of local shear and magnification distortions in the two-point function of the CMB. Optimal for current experimental noise levels, the QE will become inefficient for high-resolution, next generation survey data, such as CMB-S4 [7]. Since the primordial B-mode signal is small, it is well known that in principle one may reconstruct the lensing signal almost perfectly from the observed polarization [8], provided noise and foregrounds levels are put well below the lensing B-mode power of $\sim 5 \mu\text{K-arcmin}$. In this regime, likelihood-based methods will be able to greatly improve the lensing reconstruction by using effectively higher-order statistics of the CMB fields [8–13].

The QE power spectrum is a four-point function of the CMB maps. In addition to the sought-after lensing signal power $C_L^{\phi\phi}$, it contains other contributions, or biases, that can be characterized analytically. The dominant bias, $N_L^{(0)}$, is due to Gaussian (disconnected) contractions of the four

CMB data maps [14], with a contribution from both the instrumental noise and CMB spectra. The next-order bias (in an approach perturbative in the lensing power) is $N_L^{(1)}$. It is due to the non-Gaussian secondary trispectrum contractions of the CMB fields created by lensing, and is proportional to $C_L^{\phi\phi}$ [15]. For a perfectly Gaussian lensing map, the next-order term is $N_L^{(2)}$ which can be made negligible using suitable QE weights [16]. The QE lensing map power spectrum can then be written schematically as

$$C_L^{\hat{\phi}^{\text{QE}}\hat{\phi}^{\text{QE}}} \sim C_L^{\phi\phi} + N_L^{(0)} + N_L^{(1)}. \quad (1.1)$$

This neglects the large-scale structure and postborn bispectrum, which source an additional bias, $N_L^{(3/2)}$. As shown in Refs. [17–20], this bias is generally expected to be quite small in the QE reconstruction from polarization which we consider here. Equation (1.1) then suggests that an estimate of the lensing power can be obtained by adequate subtraction of the bias terms. For parameter inference, this is only slightly more complicated in practice, since the model dependence of the QE response and of the biases on the CMB power spectra must be taken into account for the construction of an adequate spectrum likelihood (see e.g., [4]).

A lensing map reconstructed with a likelihood-based approach is a highly complicated function of the data, and it is analytically out of reach to track the contributions to its power spectrum in a systematic manner. In this paper we investigate the lensing power spectrum from the optimal lensing map reconstruction developed in [12], showing that

*louis.legrand@unige.ch

†julien.carron@unige.ch

it shares a structure similar enough to the QE spectrum in Eq. (1.1) that it is possible to use the very same type of likelihood construction to perform unbiased parameter inference. While the biases in this case are not four-point statistics of the data, they can nevertheless be accurately obtained with very similar calculations.

Compared to recent early attempts at Monte Carlo sampling the lensing power [11,21], this provides an absolutely massive reduction in numerical cost while reaching the expected improvements of likelihood-based methods as pioneered by Ref. [8]. We demonstrate this by performing parameter inference within Λ CDM and Λ CDM + $\sum m_\nu$ models using optimal lensing map spectra as reconstructed on the full sky with curved-sky geometry. Our framework is also directly applicable to small and large sky areas inclusive of real world nonidealities like sky cuts or inhomogeneous noise. The inner workings of this first curved-sky iterative reconstruction code making this possible will be presented elsewhere [22].

State-of-the-art CMB lensing reconstructions make use of a robust, realization-dependent (RD) subtraction method of the leading noise bias, which uses QE's built partially on the data and partially on simulations [23–25]. This has two main benefits. On one hand the subtraction of the bias is more accurate, as it removes mismatch between the data and fiducial model at first order in CMB spectra. This is often of great relevance, if only because the noise properties of CMB data are in some cases only crudely under control. On the other hand, this also reduces the covariance matrix of the spectrum estimate, which otherwise typically shows large positive correlations on small scales [16,26], as well as higher variance. An important ingredient that we introduce here is a generalization of this realization-dependent bias [RD- $N_L^{(0)}$] which is suitable for the iterative estimate.

II. ITERATIVE LENSING SPECTRUM ESTIMATOR

We start by reviewing briefly the maximum *a posteriori* (MAP) lensing reconstruction and discuss its normalization. We follow the algorithm of Ref. [12], which has been demonstrated to work on data [27]. We then present our analytical predictions of the responses and biases, including RD- $N_L^{(0)}$.

A. Estimator and normalization

Let X^{dat} be an observed lensed CMB field (in what follows the Stokes parameters Q or U) including instrumental noise, and ϕ the lensing potential field (we neglect the small curl component here [9]). Assuming the unlensed CMB and noise fields are Gaussian, we can define the log-likelihood

$$\ln \mathcal{L}(X^{\text{dat}}|\phi) = -\frac{1}{2} X^{\text{dat}\dagger} \text{Cov}_\phi^{-1} X^{\text{dat}} - \frac{1}{2} \ln \det \text{Cov}_\phi, \quad (2.1)$$

where Cov_ϕ is the observed CMB (so including noise) covariance for fixed lenses, which can be modelled (to some workable approximation at least) using a fiducial model for the data noise, beam and transfer function, as well as fiducial unlensed CMB spectra with lensing and delensing operators at a given ϕ [see [12], for more details]. Using a Gaussian prior on the lensing potential with fiducial power spectrum $C_L^{\phi\phi,\text{fid}}$, we obtain the log-posterior

$$\ln \mathcal{P}(\phi|X^{\text{dat}}) = \ln \mathcal{L}(X^{\text{dat}}|\phi) - \frac{1}{2} \sum_{LM} \frac{\phi_{LM} \phi_{LM}^\dagger}{C_L^{\phi\phi,\text{fid}}}. \quad (2.2)$$

The MAP lensing estimate $\hat{\phi}^{\text{MAP}}$ is defined as the one maximizing this posterior. Since the prior gradient is proportional to ϕ we can rearrange and write¹

$$\begin{aligned} \hat{\phi}_{LM}^{\text{MAP}} &= C_L^{\phi\phi,\text{fid}} \left. \frac{2\delta \ln \mathcal{L}}{\delta \phi_{LM}^\dagger} \right|_{\phi=\hat{\phi}^{\text{MAP}}} \\ &= C_L^{\phi\phi,\text{fid}} (g_{LM}^{\text{OD}} - \langle g_{LM}^{\text{OD}} \rangle_{\hat{\phi}^{\text{MAP}}}), \end{aligned} \quad (2.3)$$

where we introduced the quadratic gradient piece g^{OD}

$$g_{LM}^{\text{OD}} = -X^{\text{dat}\dagger} \text{Cov}_{\hat{\phi}^{\text{MAP}}}^{-1} \frac{\delta \text{Cov}_{\hat{\phi}^{\text{MAP}}}}{\delta \phi_{LM}^\dagger} \text{Cov}_{\hat{\phi}^{\text{MAP}}}^{-1} X^{\text{dat}}. \quad (2.4)$$

In (2.3), its average $\langle g_{LM}^{\text{OD}} \rangle_{\hat{\phi}^{\text{MAP}}}$, which is also the first variation of the log-determinant term, is the mean field, just like in the traditional QE analysis, with the difference that the deflection field is fixed to $\hat{\phi}^{\text{MAP}}$ in the average over realizations of the fiducial model. Explicit expressions for the quadratic piece in terms of spin-weight harmonic transforms are given in Appendix A of [12].

The quadratic gradient piece has implicit dependence on $\hat{\phi}^{\text{MAP}}$ through the delensing operators involved in the inverse variance filtering step $\text{Cov}_{\hat{\phi}^{\text{MAP}}}^{-1}$, applied to the data maps. Nevertheless, Eq. (2.3) shows that the MAP estimate truly *is* quadratic in these delensed data maps.

Consider now the response to the true lenses

$$\mathcal{W}_{LML'M'} \equiv \frac{\delta \hat{\phi}_{LM}^{\text{MAP}}}{\delta \phi_{L'M'}}. \quad (2.5)$$

Since we are maximizing a posterior rather than a likelihood, we expect \mathcal{W} to be a Wiener filter; unity on resolved scales and suppressed elsewhere. Using Eq. (2.3), we can write,

$$\mathcal{W}_{LML'M'} = C_L^{\phi\phi,\text{fid}} \left[\frac{\delta g_{LM}^{\text{OD}}}{\delta X^{\text{dat}}} \frac{\delta X^{\text{dat}}}{\delta \phi} + \frac{\delta^2 \ln \mathcal{L}}{\delta \phi^\dagger \delta \phi} \frac{\delta \phi^{\text{MAP}}}{\delta \phi} \right]_{LM,L'M'}. \quad (2.6)$$

The left term of the above equation corresponds to the response of the estimator to the true lensing potential field.

¹Formally, $2\delta/\delta\phi^\dagger$ here means the harmonic transform of the variation $\delta/\delta\phi(\hat{n})$ involving only real-valued variables.

This is the direct analog to the standard QE response, for which we use similar notation \mathcal{R} . The right term comes from the implicit dependency on ϕ^{MAP} . This is equal to $-\mathcal{H}\mathcal{W}$, with \mathcal{H} the log-likelihood Hessian curvature matrix. This results in the matrix equation $\mathcal{W} = C^{\phi\phi,\text{fid}}(\mathcal{R} - \mathcal{H}\mathcal{W})$, or

$$\mathcal{W} = \left[\frac{1}{C^{\phi\phi,\text{fid}}} + \mathcal{H} \right]^{-1} \mathcal{R}. \quad (2.7)$$

In the regime where the prior is irrelevant, $1/C^{\phi\phi,\text{fid}}$ becomes negligible in front of \mathcal{H} , and assuming that the reconstruction must be unbiased, we can write $\mathcal{H}^{-1}\mathcal{R} \sim \mathbb{1}$ —the identity matrix. For practical purposes we may tentatively work with a simple isotropic approximation,

$$\mathcal{W}_L \equiv \left[\frac{1}{C_L^{\phi\phi,\text{fid}}} + \mathcal{R}_L \right]^{-1} \mathcal{R}_L = \frac{C_L^{\phi\phi,\text{fid}}}{C_L^{\phi\phi,\text{fid}} + 1/\mathcal{R}_L}. \quad (2.8)$$

This is indeed a Wiener filter with fiducial noise the inverse response $N_L = 1/\mathcal{R}_L$. This relation stands in direct analogy to the standard QE, where the reconstruction noise level $N^{(0)}$ is the exact inverse of the estimator response, provided the estimator weights are chosen to be optimal and the fiducial noise and CMB model match perfectly that of the data [28]. We stick to the notation \mathcal{R} instead of N to emphasize that it behaves like a response, independent of the data noise maps and true lenses, as we empirically confirm in the next section.

B. Likelihood construction, response and bias calculations

We use as our lensing data vector the normalized raw spectrum

$$\hat{C}_L^{\phi\phi} = C_L^{\hat{\phi}^{\text{MAP}}\hat{\phi}^{\text{MAP}}} / \mathcal{W}_L^{2,\text{fid}}, \quad (2.9)$$

where $C_L^{\hat{\phi}^{\text{MAP}}\hat{\phi}^{\text{MAP}}}$ is the pseudo power spectrum of the estimated $\hat{\phi}^{\text{MAP}}$, and $\mathcal{W}_L^{2,\text{fid}}$ is our squared fiducial Wiener filter as defined by Eq. (2.8). To build the spectrum likelihood we must consider the case where the fiducial cosmology assumed to reconstruct the lensing map, denoted by the superscript fid, is different from the sampled cosmology, denoted by θ . The negative log-likelihood is

$$-2 \ln L(\theta) = (\hat{C}_L^{\phi\phi} - C_L^{\text{pred}}(\theta)) \text{Cov}_{LL'}^{-1} (\hat{C}_L^{\phi\phi} - C_L^{\text{pred}}(\theta)), \quad (2.10)$$

where

$$C_L^{\text{pred}}(\theta) = \frac{\mathcal{W}_L^2(\theta)}{\mathcal{W}_L^{2,\text{fid}}} [C_L^{\phi\phi}(\theta) + N_L^{(1)}(\theta)] + \frac{\mathcal{W}_L^2(\theta)}{\mathcal{W}_L^{2,\text{fid}}} \frac{\mathcal{R}_L^{2,\text{fid}}}{\mathcal{R}_L^2(\theta)} \text{RD} - N_L^{(0)}, \quad (2.11)$$

and the covariance matrix $\text{Cov}_{LL'}$ is estimated from simulations (see Sec. III B).

The response \mathcal{R}_L and the $N_L^{(0)}$ and $N_L^{(1)}$ biases are obtained by assuming that they follow the same analytical expressions of the standard QE, but replacing the lensed CMB and lensing convergence spectra in the weights of the estimator by partially delensed CMB and lensing convergence spectra. We proceed as follows. Starting from fiducial and sampled lensing spectra, $C_L^{\phi\phi,\text{fid}}$ and $C_L^{\phi\phi}(\theta)$, and fiducial and sampled unlensed CMB spectra $C_\ell^{\text{EE},\text{fid}}$ and $C_\ell^{\text{EE}}(\theta)$, we iteratively compute the partially-delensed lensing spectra and the corresponding partially-lensed CMB spectra. We iterate the following three steps:

- (1) Compute the partially-lensed (at the very first iteration fully-lensed) CMB spectra for the fiducial and sampled models, together with the nonperturbative ‘grad-lensed’ spectra,² using the partially-lensed (at first fully lensed) lensing spectra.
- (2) Calculate a QE reconstruction noise level $N_L^{(0)}$, using the fiducial partially-lensed spectra as QE weights and the sampled partially-lensed spectra in the lensing response. We turn this into a cross-correlation coefficient $\rho_L = C_L^{\phi\phi,\text{fid}} / (C_L^{\phi\phi,\text{fid}} + N_L^{(0)})$ of the lensing tracer.
- (3) From ρ_L , update the fiducial and sampled partially-delensed lensing deflection spectra, given by $(1 - \rho_L^2)C_L^{\phi\phi,\text{fid}}$ and $(1 - \rho_L^2)C_L^{\phi\phi}(\theta)$. We implicitly assume that this is what the MAP reconstruction is achieving.

This procedure converges after a handful of steps. We can then calculate final estimates of the unnormalized $N^{(0)}$ and $N^{(1)}$ biases, for any choice of fiducial and sampled spectra. The response $\mathcal{R}_L(\theta)$ is calculated in the same way, using $C_\ell^{\text{EE}}(\theta)$ but always $C_L^{\phi\phi,\text{fid}}$ as lensing potential input. Our procedure is similar to the approaches of Ref. [30,31].

Using the MAP estimate and the fiducial model assumed in the reconstruction procedure, we build simulations of the observed CMB where the unlensed CMB has been deflected by $\nabla\hat{\phi}^{\text{MAP}}$, and the noise maps follow the fiducial noise model. This forms a set of simulation labeled ‘s’. We then build quadratic estimates in the same way as the

²The response function is computed with the grad-(partially) lensed spectra (defined as the cross-spectra of the CMB fields with their gradient, see Appendix C of [29]) which provide the most precise nonperturbative estimate of the response of the CMB spectra to lensing [20]. In the case of reconstruction from polarization considered here usage of these is only a tiny correction to the (partially-)lensed spectra however.

likelihood gradients g^{QD} (inclusive of the presence of $\hat{\phi}^{\text{MAP}}$ in the filters), with the difference that on one of the two legs we use the actual data map, and on the second one a simulation. This gives a set of ‘estimates’ g^{ds_i} , where by construction the response to lensing has been suppressed. We also form similar estimates $g^{s_i s_j}$ with a simulation on one leg, and another independent simulation (but still using the same deflection field $\nabla \hat{\phi}^{\text{MAP}}$) on the second leg. We then take the autopseudospectra of these two sets, noted \hat{C}_L^{di} and \hat{C}_L^{ij} respectively. The combination

$$\text{RD-}N_L^{(0)} \equiv \frac{1}{\mathcal{R}_L^2} \langle 4\hat{C}_L^{di} - 2\hat{C}_L^{ij} \rangle_{\text{MC}'s}, \quad (2.12)$$

is then our realization-dependent $N^{(0)}$ noise estimate. This $\text{RD-}N_L^{(0)}$ is used to debias the power spectrum; we only need to compute it for the final lensing estimate ϕ^{MAP} , and not at each step of the iterative lensing reconstruction.

We apply two empirical percent-level corrections, one to the Wiener filter and one to the response \mathcal{R}_L , using fiducial simulations (see below Sec. III A), and we neglect any cosmology dependence of these corrections. The parameter dependency of \mathcal{W} in the first line of (2.11) and the prefactor to $\text{RD-}N_L^{(0)}$ in the second line only enter through $C_\ell^{EE}(\theta)$, which is very tightly constrained by the CMB data, making their variation a tiny correction across a realistic MCMC chain. Furthermore, provided $\text{RD-}N_L^{(0)}$ is close to the fiducial, these two terms cancel to a very large extent. Larger is the $C_L^{\phi\phi}(\theta)$ dependence that enters $N^{(1)}$ correction, but this is also a small effect in most relevant cases. By looking at reconstructions from simulations with intentionally grossly exaggerated deviations we can sanity check our implementation. To speed up computations while sampling the likelihood, we linearize around the fiducial model for the Wiener filter, the response and the $N_L^{(1)}$ bias in $C_L^{\phi\phi}(\theta) - C_L^{\phi\phi,\text{fid}}$ and $C_\ell^{EE}(\theta) - C_\ell^{EE,\text{fid}}$, following the now standard approach of [2,4,25,32] which we found perfectly fit for our purpose.

With this linearized likelihood, we can also produce lensing-only constraints by marginalizing out the uncertainty in the true CMB spectra [4]. This introduces small additional effects from using the observed realization of the CMB spectra to calculate the filter and responses, and of augmenting the covariance matrix by $\sum_\ell \sigma_\ell^{2,E} \frac{\partial C_L^{\text{pred}}}{\partial C_\ell^{EE}} \frac{\partial C_L^{\text{pred}}}{\partial C_\ell^{EE}}$, where $\sigma^{2,E}$ is the E -spectrum covariance, assumed here to be diagonal. In our configuration this results in a 2%–4% increase on the lensing spectrum error bar at $L \sim 700$ and 1400 respectively, and less below that, so we do not consider it further in this paper.

III. RESULTS

A. Lensing power spectrum

We simulate several curved-sky, full-sky CMB realizations with variations in the ϕ and E -mode inputs, with Gaussian noise corresponding roughly to CMB-S4 wide




FIG. 1. *Upper panel:* Ratio of the empirical Wiener filter, estimated from the cross correlation of the input lensing potential ϕ^{in} and the estimated $\hat{\phi}^{\text{MAP}}$, over the fiducial Wiener filter $\mathcal{W}_L^{\text{fid}}$. Error bars show the standard error in each bin. Our analytical predictions provide a normalization of $\hat{\phi}^{\text{MAP}}$ with a bias of at most $\sim 3\%$. We use the bias estimated in this fiducial normalization to correct the \mathcal{W}_L in the rest of our analysis, as justified by the lower panel. On the largest scales the empirical Wiener filter is slightly larger than unity. This is due to our approximate treatment of the mean field which speeds up the computations but is otherwise inconsequential. *Lower panel:* Difference between the \mathcal{W}_L bias shown in the upper panel and three test cases discussed in the text: adding a large atmospheric noise component in the input map (green dots), having an input lensing spectrum lowered by 10% (red dots) compared to the fiducial one, and when we multiply this input spectrum by a wiggling function (brown dots). In these three very extreme cases the \mathcal{W}_L bias has less than 0.2% difference from the fiducial case normalization bias.

configuration with polarization noise level of $\Delta_p = \sqrt{2} \mu\text{K arcmin}$ and a beam full width at half maximum (FWHM) of 1 arcmin. We start by checking empirically the Wiener filter \mathcal{W} of the $\hat{\phi}^{\text{MAP}}$ reconstructions, all performed using the same fiducial cosmology. We estimate an effective Wiener filter from a cross spectrum to the input lensing map

$$\hat{\mathcal{W}}_L = \left\langle \frac{C_L^{\phi^{\text{in}} \hat{\phi}^{\text{MAP}}}}{C_L^{\phi^{\text{in}} \phi^{\text{in}}}} \right\rangle. \quad (3.1)$$

Figure 1 compares the empirical Wiener filter to our prediction. The lower panel of this figure confirms two important and not *a priori* obvious points making the spectrum likelihood tractable; the Wiener filter is independent of both the actual data noise and of the true lensing signal. To show this clearly we have used some extreme deviations from the fiducial model. To illustrate the first point we have included a large atmospheric noise component in the simulated data map, corresponding to the green points, given by $N_\ell = \Delta_p^2 (1 + (\frac{\ell}{700})^{-1.4})$, which is completely ignored in the covariance model of the reconstruction, using white noise power Δ_p^2 . To illustrate the second point, we have used simulations with tweaked




FIG. 2. Estimated lensing spectra subtracted by the input potential map spectrum (suppressing cosmic variance) for the QE (blue diamonds) and MAP estimators (orange diamonds). These should correspond to the sum of the $\text{RD-}N_L^{(0)}$ and $N_L^{(1)}$ biases, shown as solid lines. Also shown are the estimated spectra subtracted by the input spectrum and by $\text{RD-}N_L^{(0)}$ (in blue and orange circles for the QE and the MAP respectively). These should be dominated by the $N_L^{(1)}$ bias, with our analytical predictions shown as the dotted lines. The gray line shows the fiducial lensing spectrum.

input lensing potential power. The first has an amplitude decreased by 10% (red points), in the second a strong oscillatory signal has been superimposed to the fiducial spectrum, of the form of a factor $[1 + 0.1 \sin(2\pi\ell/200)]$ (brown points).

As already mentioned in the previous section, we correct empirically two analytical ingredients of the MAP spectrum using four independent simulations where the cosmology matches exactly the fiducial model of the reconstruction. The first correction is to the Wiener filter, as shown in Fig. 1. The second is to the response \mathcal{R} ; we rescale \mathcal{R}_L such that $\langle \text{RD-}N_L^{(0)} \rangle = 1/\mathcal{R}_L$, as should hold for a QE with optimal weights.

Figure 2 shows the estimated lensing spectra, for both QE and MAP estimators, after subtracting it by the spectrum of the input map (canceling cosmic variance), as well as by the estimated $\text{RD-}N_L^{(0)}$ bias. This shows that the estimated $\text{RD-}N_L^{(0)}$ and $N_L^{(1)}$ are accurate predictions of the biases in the estimated lensing spectra. The $N_L^{(1)}$ bias predictions starts to be less accurate at the lowest lensing multipoles, possibly in part because they are calculated in the flat-sky approximation. Since $N_L^{(1)}$ is very small on these scales this is of no significant relevance. In our likelihoods the $N_L^{(1)}$ bias for the MAP is set to zero for $L < 50$. We also see that for $L \sim 800$, the $\text{RD-}N_L^{(0)}$ bias is reduced by a factor of two with the MAP, owing to the reduction in B power achieved by the iterations. $N_L^{(1)}$, also proportional to the residual lensing power, is suppressed even further.

The robustness of our predicted lensing spectrum C_L^{pred} with respect to the fiducial and true spectra is shown in Fig. 3. We estimate $\hat{C}_L^{\phi\phi}$ for the same three extreme cases of Fig. 1. These three spectra are all normalized with the same \mathcal{W}_L




FIG. 3. Biases between the estimated and predicted spectra debiased with the fiducial $N_L^{(0)}$ (triangles) or using $\text{RD-}N_L^{(0)}$ (dots), for the same three nonfiducial, extreme realizations as on Fig. 1. The predicted spectra include the input spectrum of the map instead of the true spectrum, which cancels cosmic variance. We correct the response and $N_L^{(1)}$ bias at first order to the true lensing and polarization spectra, just as we do when sampling the cosmological model. The blue bars show our binning and one-realization statistical errors. The $\sim 1\sigma$ positive bias of the atmospheric noise case from ~ 500 to ~ 800 is due to an increased $N_L^{(1)}$ with respect to the predictions, which is well captured by including the true noise in the computation of the $N_L^{(1)}$ bias instead of the fiducial one.

corrected for the bias of the fiducial case. The predictions C_L^{pred} include either the $N_L^{(0)}$ or the $\text{RD-}N_L^{(0)}$ and the $N_L^{(1)}$ biases. For illustrative purposes, we use here the lensing spectrum of the input maps to cancel cosmic variance instead of the true lensing spectrum. In nonfiducial cases, the $N_L^{(0)}$ debiaser cannot recover an unbiased prediction of the reconstructed spectrum. Debiasing with $\text{RD-}N_L^{(0)}$ and including the corrections at first order in $C_L^{\phi\phi}(\theta) - C_L^{\phi\phi, \text{fid}}$ and $C_\ell^{EE}(\theta) - C_\ell^{EE, \text{fid}}$ for the response and $N_L^{(1)}$ bias allows to get an accurate prediction of the estimated spectrum $\hat{C}_L^{\phi\phi}$.

To obtain the covariance matrix, we estimate the QE and MAP lensing spectra from 1024 flat-sky simulations, as well as their $\text{RD-}N_L^{(0)}$. Figure 4 shows the correlation matrices (the covariance matrix normalized by the diagonal) for the QE and MAP spectra, with and without realization dependent debiasing. We see that, even without realization dependent debiasing, the MAP spectrum is less correlated between different multipole bins. The decrease of nondiagonal correlations using a realization debiaser seems to be negligible, and is only visible in the highest multipole bins ($L > 1500$). It appears this reduction of nondiagonal correlations is less important in the MAP than it is for the QE spectrum. In Fig. 5 we show the cumulative signal to noise ratio (SNR) given by

$$\text{SNR}(L_{\text{max}}) = \sqrt{\sum_{L_{\text{min}}}^{L_{\text{max}}} C_L^{\phi\phi, \text{fid}} \text{Cov}_{LL'}^{-1} C_{L'}^{\phi\phi, \text{fid}}}. \quad (3.2)$$




FIG. 4. Correlation matrices of MAP (top row) and QE (bottom row) lensing convergence spectra from 1024 flat-sky simulations, binned with a step $\Delta L = 100$ between $L_{\min} = 10$ and $L_{\max} = 2000$. In the right panels the spectra are debiased with $\text{RD-}N_L^{(0)}$ while in the left panels they are not.

The cumulative SNR of the QE with a realization dependent debiaser saturates around 190 at $L_{\max} = 1000$, while for the MAP the SNR keeps increasing up to ~ 340 at $L_{\max} = 1500$. The information gained with the use of the $\text{RD-}N_L^{(0)}$ is less important for the MAP than for the QE, because the covariance matrix of the MAP is already optimal in the range of scales which contain most of the signal.

B. Cosmological parameter estimation

We demonstrate the potential of our MAP spectrum estimator to obtain constraints on cosmological models for a CMB-S4 experiment. First, using only our lensing likelihood we measure the constraints on the $\sigma_8 \Omega_m^{0.27}$ parameter combination. Second, we let free the sum of neutrino masses and obtain constraints on it by adding CMB and baryonic acoustic oscillations (BAO) likelihoods. To check that our pipeline is unbiased, we simulate two datasets, each in a different cosmology. The first one is the *Planck* FFP10 model,³ with $\sum m_\nu = 0.06$ eV. This is our fiducial cosmology for all lensing reconstructions. The second one has an higher neutrino total mass of $\sum m_\nu = 0.1$ eV, with

³https://github.com/carronj/plancklens/blob/master/plancklens/data/cls/FFP10_wdipole_params.ini.


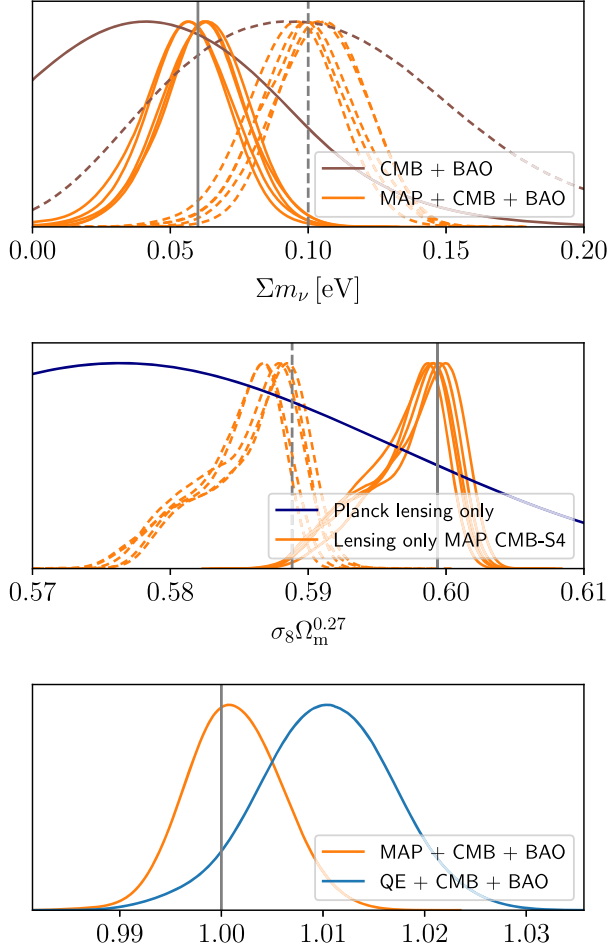


FIG. 5. Signal to noise ratio for the MAP (orange lines) and the QE (blue lines) lensing power spectra. For dashed lines the spectra are debiased with $\text{RD-}N_L^{(0)}$, while plain lines are not. The MAP increases the SNR by 80% compared to the QE, from 190 to 340.

three massive neutrinos in a normal hierarchy and two degenerate mass eigenstates, while keeping the angular size of the sound horizon fixed. The latter is performed by changing the Hubble constant slightly. For both cosmologies we simulate five full-sky lensed CMB, with CMB-S4 noise level, and we estimate the MAP lensing spectrum from the polarization maps. These spectra are binned between $L_{\min} = 10$ and $L_{\max} = 1500$ with a step $\Delta L = 10$ up to $L = 100$, and a step $\Delta L = 50$ above. The covariance matrix for $L > 100$ is obtained from lensing spectra estimated with 1024 flat-sky simulations of 645 deg^2 each. This matrix is normalized in order to reproduce a 40% sky fraction. The larger scales, $L < 100$, use an analytical Gaussian covariance with same sky fraction. We sample our likelihoods with the code Cobaya [33], using an adaptive MCMC sampler [34,35], and the Boltzmann code CAMB [36,37].

The sensitivity of CMB lensing to Λ CDM cosmological parameters is discussed in [38–40]. There is a three parameter degeneracy $\sigma_8\text{-}\Omega_m\text{-}H_0$ ‘tube’, which projects onto a tightly constrained $\sim \sigma_8 \Omega_m^{0.25}$. For our lensing-only constraints, we use the same priors as the *Planck* analysis [4], most notably a prior on the baryon density from abundance measurements that constrains the sound horizon (a prior much weaker than the constraints expected from CMB-S4). The marginalized posterior on the CMB lensing parameter is shown in the lower panel of Fig. 6, where we also show for comparison the constraints from the *Planck* lensing-only analysis [4]. For both input cosmologies we recover an unbiased estimate of the $\sigma_8 \Omega_m^{0.27}$ parameter combination, with constraints about seven times better than current best data (the 0.27 exponent was found with a principal component analysis of our chains).

CMB lensing is sensitive to the sum of neutrino masses through the suppression of the growth on scales smaller than the free streaming scale [41]. Combining primordial CMB spectra and BAO can break the CMB degeneracies by putting constraints on the sound horizon at low and high redshift [42]. Good knowledge of the optical depth to reionization τ fixes then the primordial fluctuation



$$I = (1 + \sum m_\nu - (\sum m_\nu)^{\text{fid}}) (\Omega_m / \Omega_m^{\text{fid}})^{-1.7} (\tau / \tau^{\text{fid}})^{-0.18}$$

FIG. 6. *Upper panel:* Λ CDM + $\sum_\nu m_\nu$ marginalized posterior on the sum of neutrino masses for our test CMB-S4 lensing likelihoods. Constraints from combined CMB-S4 and DESI BAO are shown in brown, while constraints including further the MAP lensing spectrum are shown in orange. For solid lines the true neutrino mass is of 0.06 eV (solid gray line), corresponding to the fiducial cosmology in our lensing reconstruction. For dashed lines, the true neutrino mass is of 0.1 eV (dashed gray line) while the fiducial cosmology for the reconstruction stays the same. For each dataset we show the posterior of five different lensed CMB realizations. *Medium panel:* The orange lines show Λ CDM marginalized constraints on the CMB-lensing parameter $\sigma_8 \Omega_m^{0.27}$ with our MAP lensing-only likelihood for a CMB-S4 survey. The dark blue line shows the posterior from the latest *Planck* data for comparison [4]. The solid, and dashed lines stands for the two inputs cosmologies, as above. The shape of the posterior is driven by the remaining correlations in the $\sigma_8 - \Omega_m$ plane. *Lower panel:* Λ CDM + $\sum_\nu m_\nu$ marginalized posterior on the derived parameter I , for the combination of our lensing, CMB-S4 and DESI BAO likelihoods, in the fiducial dataset. The blue line uses a QE to estimate the lensing spectrum, while the orange line uses the MAP estimator. On this combined parameter the MAP spectrum is able to reduce the marginalized 1σ constraint by 30%.

amplitude. Combining CMB + BAO + CMB lensing then provides a constraint on the neutrinos total mass. Current tightest constraints on the neutrino masses obtained by combining CMB + BAO + CMB lensing datasets are of $\sum m_\nu < 0.12$ eV (95% confidence level) [43], and the CMB-S4 + DESI BAO + CMB-S4 lensing combination is expected to be able to detect to high significance the minimal neutrino mass allowed by terrestrial experiments [7]. We sample the seven parameters of the Λ CDM + $\sum m_\nu$ model for the combined posterior including CMB and BAO likelihoods, with or without our MAP lensing likelihood. We include a Gaussian prior $\sigma_\tau = 0.002$. This corresponds roughly to the cosmic variance limit for a full sky polarization survey such as LiteBIRD [44]. Our simulated CMB data vector is a set of simulated TT, TE, and EE CMB unlensed spectra ranging from $\ell_{\min} = 30$ to $\ell_{\max} = 3000$, with covariance rescaled to 40% of the sky observations. The BAO likelihood reproduces the one used to forecast the DESI survey [45], following the recipe of [46]. We include the main galaxy sample, the low redshift bright galaxy sample, and the high redshift Lyman- α quasar survey, for a total redshift range from $z = 0.05$ to $z = 3.55$. The marginalized constraints on the sum of neutrino masses are shown in the upper panel of Fig. 6. In both the low- and high-mass cosmologies, our likelihood pipeline obtains an unbiased estimate of the true neutrino masses. We obtain a 3.8σ (respectively 6.3σ in the high mass case) detection of massive neutrinos, consistent with the forecasts presented in the CMB-S4 science book [7].

While the total signal to noise ratio of the lensing spectrum with the MAP is 80% higher than with the QE, the marginalized constraint on the sum of neutrino mass is only slightly improved. We obtain $\sigma_{M_\nu} = 0.016$ eV and $\sigma_{M_\nu} = 0.017$ eV for the MAP and QE respectively. This rather small difference is consistent with the Fisher analysis of [47]. This probably means that the marginalized constraint is not driven by the improvement in statistical power brought by the MAP estimator, but by degeneracies between cosmological parameters. To test this, we perform a principal component analysis of our chains on the parameters $\sum m_\nu$, Ω_m and τ , among the main parameters impacting the amplitude of the lensing power spectrum. We found that the combined parameter $I = (1 + \sum m_\nu - (\sum m_\nu)^{\text{fid}}) (\Omega_m / \Omega_m^{\text{fid}})^{-1.7} (\tau / \tau^{\text{fid}})^{-0.18}$, gets better constraints with the MAP than with the QE, as shown on the lower panel of Fig. 6. The 1σ constraints are of $\sigma_I = 6.7$ meV and $\sigma_I = 4.7$ meV with the QE and MAP respectively, corresponding to an improvement of 30%. This does not yet match the statistical improvement obtained by the MAP reconstruction, so there might still be some degeneracies left.

In order to assess the impact of the reionization optical depth prior on the neutrino mass constraints, we compute




FIG. 7. *Upper panel:* Marginalized 1σ Fisher constraints on the sum of neutrino masses as a function of standard deviation of the Gaussian prior on τ . Blue (orange) line combines the CMB-S4 and DESI BAO likelihoods with the QE (resp. MAP) lensing likelihood. The dashed line shows the prior of $\sigma(\tau) = 0.002$ we used in our MCMC analysis. *Lower panel:* Improvement (in percent) on the neutrino mass constraint when using the MAP instead of the QE lensing spectrum.

the Fisher matrix for the combination of our CMB-S4 lensing potential, CMB-S4 unlensed and DESI BAO likelihoods as above. We use either the QE or the MAP lensing likelihoods. Figure 7 shows the marginalized 1σ constraints on $\sum m_\nu$ as a function of the standard deviation of the Gaussian prior on τ . When using the MAP lensing likelihood, we see that the constraint is of $\sigma_{M_\nu} = 0.016$ eV when τ is fixed [i.e., $\sigma(\tau)_{\text{prior}} = 0$] and increases to $\sigma_{M_\nu} = 0.027$ eV with a looser prior on τ . This is in agreement with our constraints from the MCMC chains above [using the prior of $\sigma(\tau)_{\text{prior}} = 0.002$], and is also in agreement with the CMB-S4 forecasts [7]. Here also we see that the improvement from the QE to the MAP constraints is of only a few percents, and at best of 5% when τ is fixed. As we showed above, there are still degeneracies between cosmological parameters which have to be broken to reach the full statistical power of the MAP lensing reconstruction.

IV. DISCUSSION

In this paper we introduced a new CMB lensing power spectrum estimator, using optimal likelihood-based lensing reconstruction which carries much more signal to noise for experiments with low polarization noise. Keeping the numerical cost under control, in principle directly applicable to masked data or with other nonidealities, with a robust realization-dependent debiaser, this essentially solves several practical challenges facing reconstruction of the lensing power spectrum beyond the QE.

Our spectrum estimator uses altogether a single MAP lensing reconstruction, performed assuming a fiducial cosmological model. While the MAP reconstruction is much more expensive than its QE counterpart and dominates the overall numerical cost, the entire pipeline remains all things considered quite economical. As a point of comparison to the literature, the recent iterative

spectrum estimation proposal of Ref. [48] performs of $O(10^3)$ MAP reconstructions. Owing to the somewhat arbitrary choice of fiducial cosmology, our estimator will be somewhat sub-optimal. However, our cosmological model is so tightly constrained already at the present time, that this can hardly be more than a percent-level effect.

We neglected several issues. CMB foregrounds (such as Galactic dust emission, Sunyaev-Zel'dovich effect, radio sources, or the cosmic infrared background), particularly in temperature, can bias the lensing reconstruction by creating non-Gaussian signatures in the observed CMB [49–57]. It is not known yet what is the size of these biases in the MAP reconstruction. We also neglected the bias from the large-scale bispectrum [17–20], more relevant on the smaller scales probed by CMB-S4 lensing than it is today. This bias is especially important when performing a tomographic analysis of CMB lensing in cross-correlation with Galaxy surveys [58–62], a promising probe of the growth of structures. Here also it is not yet known what is the importance of this bias when using the MAP reconstruction. For these cross-correlations, our results will be useful to model accurately the normalization of the MAP estimate. Finally, we did not discuss the origin of the small corrections to the fiducial Wiener filter and response. As we showed, this does not limit nor bias our analysis, but there is room for a more precise analytical understanding of our results.

We did not consider masking and other instrumental nonidealities either in this work. The MAP reconstruction is known to work on realistic data. It is worth noting that all ingredients introduced here, inclusive of $\text{RD-}N_L^{(0)}$, can be obtained in just the same way on a masked sky. It is well known that close the mask boundaries, the QE isotropic normalization is inaccurate, and must be corrected for by small Monte Carlo correction. In a preliminary analysis, we found as expected the same behavior for the MAP reconstruction, with a bias less than a few percents. This bias is of similar size to the one for the full-sky analysis presented in this paper, and hence it is simple to account for them jointly. Another important difference in the masked case is the much larger size of the low- L mean field, induced by the large-scale anisotropies from notably the mask and scanning pattern. An exact treatment of this mean field would in principle require the analysis of a number of simulations at each iteration step, multiplying the numerical cost by this number. However, Ref. [27] already demonstrated that one can make profit of the small size of the $\hat{\phi}$ dependence of the mean field to use the QE mean-field estimate in each of the MAP iterations. For this reason we do not expect the mean field to slow down the pipeline. We leave for future work the full adaptation of this analysis pipeline to more realistic simulations of the CMB-S4 survey (or other planned observations), or actual data.

ACKNOWLEDGMENTS

The authors wish to thank Giulio Fabbian, Antony Lewis, Blake Sherwin, and Kimmy Wu for useful discussions and comments, as well as the anonymous

referees to the first version of this paper. The authors acknowledge support from a SNSF Eccellenza Professorial Fellowship (No. 186879).

-
- [1] Antony Lewis and Anthony Challinor, Weak gravitational lensing of the CMB, *Phys. Rep.* **429**, 1 (2006).
- [2] Blake D. Sherwin *et al.*, Two-season Atacama Cosmology Telescope polarimeter lensing power spectrum, *Phys. Rev. D* **95**, 123529 (2017).
- [3] Y. Omori *et al.*, A 2500 deg² CMB lensing map from combined South Pole Telescope and Planck data, *Astrophys. J.* **849**, 124 (2017).
- [4] N. Aghanim *et al.* (Planck Collaboration), Planck 2018 results. VIII. Gravitational lensing, *Astron. Astrophys.* **641**, A8 (2020).
- [5] 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).
- [6] Takemi Okamoto and Wayne Hu, CMB lensing reconstruction on the full sky, *Phys. Rev. D* **67**, 083002 (2003).
- [7] Kevork N. Abazajian *et al.* (CMB-S4 Collaboration), CMB-S4 Science Book, First Edition, [arXiv:1610.02743](https://arxiv.org/abs/1610.02743).
- [8] Christopher M. Hirata and Uros Seljak, Reconstruction of lensing from the cosmic microwave background polarization, *Phys. Rev. D* **68**, 083002 (2003).
- [9] Christopher M. Hirata and Uros Seljak, Analyzing weak lensing of the cosmic microwave background using the likelihood function, *Phys. Rev. D* **67**, 043001 (2003).
- [10] Marius Millea, Ethan Anderes, and Benjamin D. Wandelt, Bayesian delensing of CMB temperature and polarization, *Phys. Rev. D* **100**, 023509 (2019).
- [11] Marius Millea, Ethan Anderes, and Benjamin D. Wandelt, Sampling-based inference of the primordial CMB and gravitational lensing, *Phys. Rev. D* **102**, 123542 (2020).
- [12] Julien Carron and Antony Lewis, Maximum a posteriori CMB lensing reconstruction, *Phys. Rev. D* **96**, 063510 (2017).
- [13] Julien Carron, Optimal constraints on primordial gravitational waves from the lensed CMB, *Phys. Rev. D* **99**, 043518 (2019).
- [14] Wayne Hu and Takemi Okamoto, Mass reconstruction with CMB polarization, *Astrophys. J.* **574**, 566 (2002).
- [15] Michael H. Kesden, Asantha Cooray, and Marc Kamionkowski, Lensing reconstruction with CMB temperature and polarization, *Phys. Rev. D* **67**, 123507 (2003).
- [16] Duncan Hanson, Anthony Challinor, George Efstathiou, and Pawel Bielewicz, CMB temperature lensing power reconstruction, *Phys. Rev. D* **83**, 043005 (2011).
- [17] Vanessa Böhm, Marcel Schmittfull, and Blake D. Sherwin, Bias to CMB lensing measurements from the bispectrum of large-scale structure, *Phys. Rev. D* **94**, 043519 (2016).
- [18] Vanessa Böhm, Blake D. Sherwin, Jia Liu, J. Colin Hill, Marcel Schmittfull, and Toshiya Namikawa, Effect of non-Gaussian lensing deflections on CMB lensing measurements, *Phys. Rev. D* **98**, 123510 (2018).
- [19] Dominic Beck, Giulio Fabbian, and Josquin Errard, Lensing reconstruction in post-born cosmic microwave background weak lensing, *Phys. Rev. D* **98**, 043512 (2018).
- [20] Giulio Fabbian, Antony Lewis, and Dominic Beck, CMB lensing reconstruction biases in cross-correlation with large-scale structure probes, *J. Cosmol. Astropart. Phys.* **10** (2019) 057.
- [21] M. Millea *et al.*, Optimal cosmic microwave background lensing reconstruction and parameter estimation with SPTpol data, *Astrophys. J.* **922**, 259 (2021).
- [22] The CMB-S4 Collaboration, CMB-S4: Iterative internal delensing and r constraints (to be published).
- [23] Toshiya Namikawa, Duncan Hanson, and Ryuichi Takahashi, Bias-Hardened CMB lensing, *Mon. Not. R. Astron. Soc.* **431**, 609 (2013).
- [24] K. T. Story *et al.* (SPT Collaboration), A measurement of the cosmic microwave background gravitational lensing potential from 100 square degrees of SPTpol data, *Astrophys. J.* **810**, 50 (2015).
- [25] P. A. R. Ade *et al.* (Planck Collaboration), Planck 2015 results. XV. Gravitational lensing, *Astron. Astrophys.* **594**, A15 (2016).
- [26] Julien Peloton, Marcel Schmittfull, Antony Lewis, Julien Carron, and Oliver Zahn, Full covariance of CMB and lensing reconstruction power spectra, *Phys. Rev. D* **95**, 043508 (2017).
- [27] S. Adachi *et al.* (POLARBEAR Collaboration), Internal Delensing of Cosmic Microwave Background polarization B -Modes with the POLARBEAR Experiment, *Phys. Rev. Lett.* **124**, 131301 (2020).
- [28] Abhishek S. Maniyar, Yacine Ali-Haïmoud, Julien Carron, Antony Lewis, and Mathew S. Madhavacheril, Quadratic estimators for CMB weak lensing, *Phys. Rev. D* **103**, 083524 (2021).
- [29] Antony Lewis, Anthony Challinor, and Duncan Hanson, The shape of the CMB lensing bispectrum, *J. Cosmol. Astropart. Phys.* **03** (2011) 018.
- [30] Kendrick M. Smith, Duncan Hanson, Marilena LoVerde, Christopher M. Hirata, and Oliver Zahn, Delensing CMB polarization with external datasets, *J. Cosmol. Astropart. Phys.* **06** (2012) 014.
- [31] Selim C. Hotinli, Joel Meyers, Cynthia Trendafilova, Daniel Green, and Alexander van Engelen, The benefits of CMB delensing, *J. Cosmol. Astropart. Phys.* **04** (2022) 020.

- [32] G. Simard *et al.* (SPT Collaboration), Constraints on cosmological parameters from the angular power spectrum of a combined 2500 deg² SPT-SZ and Planck gravitational lensing map, *Astrophys. J.* **860**, 137 (2018).
- [33] Jesus Torrado and Antony Lewis, Cobaya: Code for Bayesian analysis of hierarchical physical models, *J. Cosmol. Astropart. Phys.* **05** (2021) 057.
- [34] Antony Lewis and Sarah Bridle, Cosmological parameters from CMB and other data: A Monte Carlo approach, *Phys. Rev. D* **66**, 103511 (2002).
- [35] Antony Lewis, Efficient sampling of fast and slow cosmological parameters, *Phys. Rev. D* **87**, 103529 (2013).
- [36] Antony Lewis, Anthony Challinor, and Anthony Lasenby, Efficient computation of CMB anisotropies in closed FRW models, *Astrophys. J.* **538**, 473 (2000).
- [37] Cullan Howlett, Antony Lewis, Alex Hall, and Anthony Challinor, CMB power spectrum parameter degeneracies in the era of precision cosmology, *J. Cosmol. Astropart. Phys.* **04** (2012) 027.
- [38] Z. Pan, L. Knox, and M. White, Dependence of the cosmic microwave background lensing power spectrum on the matter density, *Mon. Not. R. Astron. Soc.* **445**, 2941 (2014).
- [39] P. A. R. Ade *et al.* (Planck Collaboration), Planck 2015 results. XV. Gravitational lensing, *Astron. Astrophys.* **594**, A15 (2016).
- [40] F. Bianchini *et al.* (SPT Collaboration), Constraints on cosmological parameters from the 500 deg² SPTpol lensing power spectrum, *Astrophys. J.* **888**, 119 (2020).
- [41] Julien Lesgourgues and Sergio Pastor, Massive neutrinos and cosmology, *Phys. Rep.* **429**, 307 (2006).
- [42] Will J. Percival *et al.* (2dFGRS Team Collaboration), Parameter constraints for flat cosmologies from CMB and 2dFGRS power spectra, *Mon. Not. R. Astron. Soc.* **337**, 1068 (2002).
- [43] N. Aghanim *et al.* (Planck Collaboration), Planck 2018 results. VI. Cosmological parameters, *Astron. Astrophys.* **641**, A6 (2020); **652**, C4(E) (2021).
- [44] M. Hazumi *et al.* (LiteBIRD Collaboration), LiteBIRD: JAXA's new strategic L-class mission for all-sky surveys of cosmic microwave background polarization, *Proc. SPIE Int. Soc. Opt. Eng.* **11443**, 114432F (2020).
- [45] Amir Aghamousa *et al.* (DESI Collaboration), The DESI experiment Part I: Science, targeting, and survey design, [arXiv:1611.00036](https://arxiv.org/abs/1611.00036).
- [46] Andreu Font-Ribera, Patrick McDonald, Nick Mostek, Beth A. Reid, Hee-Jong Seo, and An Slosar, DESI and other dark energy experiments in the era of neutrino mass measurements, *J. Cosmol. Astropart. Phys.* **05** (2014) 023.
- [47] R. Allison, P. Caucal, E. Calabrese, J. Dunkley, and T. Louis, Towards a cosmological neutrino mass detection, *Phys. Rev. D* **92**, 123535 (2015).
- [48] Marius Millea and Uros Seljak, MUSE: Marginal unbiased score expansion and application to CMB lensing, *Phys. Rev. D* **105**, 103531 (2022).
- [49] A. van Engelen, S. Bhattacharya, N. Sehgal, G. P. Holder, O. Zahn, and D. Nagai, CMB lensing power spectrum biases from galaxies and clusters using high-angular resolution temperature maps, *Astrophys. J.* **786**, 13 (2014).
- [50] Stephen J. Osborne, Duncan Hanson, and Olivier Doré, Extragalactic foreground contamination in temperature-based CMB lens reconstruction, *J. Cosmol. Astropart. Phys.* **03** (2014) 024.
- [51] Simone Ferraro and J. Colin Hill, Bias to CMB lensing reconstruction from temperature anisotropies due to large-scale galaxy motions, *Phys. Rev. D* **97**, 023512 (2018).
- [52] Mathew S. Madhavacheril and J. Colin Hill, Mitigating foreground biases in CMB lensing reconstruction using cleaned gradients, *Phys. Rev. D* **98**, 023534 (2018).
- [53] Emmanuel Schaan and Simone Ferraro, Foreground-Immune Cosmic Microwave Background Lensing with Shear-Only Reconstruction, *Phys. Rev. Lett.* **122**, 181301 (2019).
- [54] Nishant Mishra and Emmanuel Schaan, Bias to CMB lensing from lensed foregrounds, *Phys. Rev. D* **100**, 123504 (2019).
- [55] Noah Sailer, Emmanuel Schaan, and Simone Ferraro, Lower bias, lower noise CMB lensing with foreground-hardened estimators, *Phys. Rev. D* **102**, 063517 (2020).
- [56] Omar Darwish, Blake D. Sherwin, Noah Sailer, Emmanuel Schaan, and Simone Ferraro, Optimizing foreground mitigation for CMB lensing with combined multifrequency and geometric methods, [arXiv:2111.00462](https://arxiv.org/abs/2111.00462).
- [57] Margherita Lembo, Giulio Fabbian, Julien Carron, and Antony Lewis, CMB lensing reconstruction biases from masking extragalactic sources, [arXiv:2109.13911](https://arxiv.org/abs/2109.13911) [Phys. Rev. D (to be published)].
- [58] Christopher M. Hirata, Shirley Ho, Nikhil Padmanabhan, Uros Seljak, and Neta A. Bahcall, Correlation of CMB with large-scale structure: II. Weak lensing, *Phys. Rev. D* **78**, 043520 (2008).
- [59] Emmanuel Schaan, Elisabeth Krause, Tim Eifler, Olivier Doré, Hironao Miyatake, Jason Rhodes, and David 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).
- [60] Marcel Schmittfull and Uros Seljak, Parameter constraints from cross-correlation of CMB lensing with galaxy clustering, *Phys. Rev. D* **97**, 123540 (2018).
- [61] T. M. C. Abbott *et al.* (DES, SPT Collaborations), Dark Energy Survey Year 1 results: Joint analysis of galaxy clustering, galaxy lensing, and CMB lensing two-point functions, *Phys. Rev. D* **100**, 023541 (2019).
- [62] S. Ilić *et al.* (Euclid Collaboration), *Euclid* preparation: XV. Forecasting cosmological constraints for the *Euclid* and CMB joint analysis, *Astron. Astrophys.* **657**, A91 (2022).