

An iterative CMB lensing estimator minimizing instrumental noise bias

Louis Legrand^{1,2,*} Blake Sherwin,^{1,2} Anthony Challinor^{3,1,2} Julien Carron⁴ and Gerrit S. Farren^{5,6}

¹*Department of Applied Mathematics and Theoretical Physics,*

University of Cambridge, Wilberforce Road, Cambridge CB3 0WA, United Kingdom

²*Kavli Institute for Cosmology, Cambridge, Madingley Road, Cambridge CB3 0HA, United Kingdom*

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

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

⁵*Physics Division, Lawrence Berkeley National Laboratory,*

1 Cyclotron Road, Berkeley, California 94720, USA

⁶*Berkeley Center for Cosmological Physics, University of California, Berkeley, California 94720, USA*



(Received 18 July 2025; accepted 29 October 2025; published 17 November 2025)

Noise maps from cosmic microwave background (CMB) experiments are generally statistically anisotropic, due to scanning strategies, atmospheric conditions, or instrumental effects. Any mismodeling of this complex noise can bias the reconstruction of the lensing potential and the measurement of the lensing power spectrum from the observed CMB maps. We introduce a new CMB lensing estimator based on the maximum *a posteriori* (MAP) reconstruction that is minimally sensitive to these instrumental noise biases. By modifying the likelihood to rely exclusively on correlations between CMB map splits with independent noise realizations, we minimize autocorrelations that contribute to biases. In the regime of many independent splits, this maximum closely approximates the optimal MAP reconstruction of the lensing potential. In simulations, we demonstrate that this method is able to determine lensing observables that are immune to any noise mismodeling with a negligible cost in signal-to-noise ratio. Our estimator enables unbiased and nearly optimal lensing reconstruction for next-generation CMB surveys.

DOI: [10.1103/8kt3-d8z6](https://doi.org/10.1103/8kt3-d8z6)

I. INTRODUCTION

The gravitational lensing of the cosmic microwave background (CMB) is a powerful probe of the large-scale structure of the Universe. It allows us to reconstruct the projected mass distribution along the line of sight and to constrain the clustering of matter, the sum of neutrino masses, and the properties of dark energy [1].

Quadratic estimators (QEs), which optimally combine weighted pairs of CMB modes, are typically used to reconstruct the lensing potential from the statistical anisotropies it creates in the CMB [2]. However, it has been established that the QE approach is suboptimal for low-noise CMB observations, in particular with deep polarization surveys (e.g., [3–5]). While the QE is equivalent to a Gaussian approximation of the lensing potential likelihood, maximum *a posteriori* (MAP) estimators use all of the information in the CMB maps to reconstruct the lensing

potential. For a CMB-S4 configuration, the MAP estimator is expected to halve the errors on the reconstruction noise power spectrum over a wide range of scales, as compared to the QE [6].

CMB lensing estimators are sensitive to the modeling of the noise and other sources of statistical anisotropies in the CMB maps. This shows up in two ways. Firstly, any nonlensing source of anisotropies in the CMB maps, such as masking, scanning strategies, and atmospheric effects, can lead to a bias in the reconstructed lensing potential map, known as the *mean field*. Secondly, for the QE, the power spectrum of the reconstructed lensing map is a four-point function of the CMB maps, so it includes the lensing signal we want to measure, plus other four-point contractions. The dominant term is the disconnected four-point function of the Gaussian CMB and noise [giving rise to what is referred to as the $N_L^{(0)}$ bias], which is present even in the absence of gravitational lensing. At small scales, this bias can be orders of magnitude larger than the signal. Thus, in order to avoid biases in the lensing map and in the lensing power spectrum, the mean field and $N_L^{(0)}$ bias are estimated with Monte Carlo simulations, reproducing the noise pattern and scanning strategy of the CMB experiment [7–10], and subtracted.

*Contact author: ll783@cam.ac.uk

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

However, accurate simulations of the observed CMB are difficult to achieve, in particular for ground-based surveys with complex atmospheric effects [11]. One solution introduced in [12] is to take advantage of the fact that the atmospheric and instrumental noise in the CMB maps vary with short correlation times. This means that if we are able to split the CMB maps into several sets observed at different times, they will have independent atmospheric and instrumental noise, while the CMB component (and potential Galactic and extragalactic foregrounds) will be stable. This *cross-only* QE uses different pairs of CMB maps in each leg of the estimator and combines all possible pairs, avoiding autocorrelations in the lensing power spectrum. Because the noise in each leg of the QE lensing power spectrum is independent, this estimator effectively nulls the noise component in the mean field and in the $N_L^{(0)}$ biases. Reference [12] demonstrated that in the limit of a very large number of splits, or when the CMB noise level is much lower than the signal, the lensing power spectrum can be reconstructed with a negligible cost in signal-to-noise ratio. This allows for an unbiased estimate of the CMB lensing power spectrum, without the need for accurate simulations of the various sources of noise.

Alternatively, Ref. [13] proposed a method to remove the disconnected biases completely, i.e., from both noise and signal, by applying on the CMB maps mutually exclusive annular filters in Fourier space. This method was extended in [14], relying on the fact that the Gaussian noise bias comes from a relatively small number of specific multipole configurations of the CMB four-point function, which can be removed from the lensing power spectrum. However, these two methods break for highly anisotropic noise configurations, such as the ones we are interested in for our work.

In this paper we extend the cross-only estimator of Ref. [12] to the MAP lensing estimator. The main difficulty is that, by definition, the MAP estimator uses all the information in the CMB maps to reconstruct the lensing potential. So, if we split the observed CMB into several maps observed at different times, the likelihood formalism will leverage both auto- and cross-correlations between these different maps to reconstruct the lensing signal. We will develop a new method that allows us to neglect the autocorrelation parts in the likelihood maximization algorithm. In practice, this comes down to replacing the standard likelihood estimator by a loss function that mimics the likelihood while neglecting the autocorrelations of the split maps. While this loss function has in principle no guarantee to converge to the maximum of the likelihood, we show that in scenarios where the CMB signal dominates over the noise, the maximization of the loss function converges to a lensing map very close to optimal, while greatly reducing the noise contribution in the mean field and in the biases of the lensing power spectrum, at a negligible cost in the signal-to-noise ratio.

The paper is organized as follows. In Sec. II we review the formalism of the MAP lensing estimator and introduce the cross-only MAP estimator. In Sec. III we validate the estimator on simulations for different CMB configurations. We conclude in Sec. IV. Finally, we derive the mean field of our cross-only estimator in the Appendix.

II. LENSING ESTIMATORS

A. Maximum *a posteriori* lensing estimator

We review here the formalism of the MAP lensing estimator, as introduced in [4,15]. We work here in the flat-sky formalism for simplicity, but the approach naturally extends to the curved sky. Note that in Sec. III, we will work with full-sky simulations and use curved-sky estimators.

We model the observed CMB maps as the data vector

$$X^{\text{dat}} = \mathcal{B}\mathcal{D}_\alpha \tilde{X} + n. \quad (2.1)$$

Here, the column vector $X^{\text{dat}}(\hat{n}) = (T^{\text{dat}}, Q^{\text{dat}}, U^{\text{dat}})^T(\hat{n})$ is the observed CMB data vector for temperature and Q and U Stokes parameters in real space, $\tilde{X}_l = (T_l, E_l, B_l)^T$ are the unlensed CMB fields in multipole space (with E_l and B_l the E - and B -mode polarization multipoles, respectively), $\mathcal{B}$ is the beam and transfer function of the instrument, and n is the noise in the maps. The lensing deflection field α is related to the lensing potential with $\alpha = \nabla\phi$ (we neglect the curl potential), and $\mathcal{D}_\alpha$ is the map-synthesis and deflection operator that transforms the unlensed CMB modes in multipole space into the lensed Stokes parameters in real space.

The covariance of this data vector at a fixed lensing deflection field α is

$$\text{Cov}_\alpha = \langle X^{\text{dat}} X^{\text{dat},\dagger} \rangle_\alpha = \mathcal{B}\mathcal{D}_\alpha C^{\text{unl}} \mathcal{D}_\alpha^\dagger \mathcal{B}^\dagger + N, \quad (2.2)$$

where the covariance C^{unl} contains the unlensed CMB power spectra and N is the noise covariance matrix. The delensing operation is $\mathcal{D}_\alpha^\dagger$, and an efficient implementation of this operator is discussed in [16].

At fixed lensing deflections, the CMB follows a Gaussian likelihood $\mathcal{L}$. Assuming a Gaussian prior, $-2\ln \mathcal{P}(\alpha) = \sum_L |\phi_L|^2 / C_L^{\phi\phi} + \text{const}$, on the lensing field, we can write, up to a constant, the log-posterior of the deflection field as

$$\begin{aligned} \ln \mathcal{P}(\alpha | X^{\text{dat}}) &= \ln \mathcal{L}(X^{\text{dat}} | \alpha) + \ln \mathcal{P}(\alpha) \\ &= -\frac{1}{2} X^{\text{dat},T} \text{Cov}_\alpha^{-1} X^{\text{dat}} - \frac{1}{2} \ln \det \text{Cov}_\alpha \\ &\quad - \frac{1}{2} \sum_L \frac{|\phi_L|^2}{C_L^{\phi\phi}}. \end{aligned} \quad (2.3)$$

Following the procedure of [4,15], the maximum *a posteriori* deflection field α^{MAP} is found by Newton-Raphson

iterations on this posterior. This involves computing the gradient and the curvature of the log-posterior with respect to the deflection field. The iterations are converged when the gradient of the log-posterior vanishes.

The gradient of the term which is quadratic in the data can be written as¹

$$g_a^{\text{OD}} = [\bar{X}]^x [\mathcal{D}_\alpha \nabla_a X^{\text{WF}}]_x, \quad (2.4)$$

where the index a denotes one of the two Cartesian coordinates on the flat sky, $a \in (1, 2)$, and the brackets are for summing over the stokes parameters $x \in (T, Q, U)$, and we have introduced the filtered maps

$$\begin{aligned} \bar{X} &\equiv \mathcal{B}^\dagger \text{Cov}_\alpha^{-1} X^{\text{dat}}, \\ X^{\text{WF}} &\equiv C^{\text{unl}} \mathcal{D}_\alpha^\dagger \mathcal{B}^\dagger \text{Cov}_\alpha^{-1} X^{\text{dat}}. \end{aligned} \quad (2.5)$$

Note that the Wiener-filtered field X^{WF} is defined in harmonic space and the notation $\nabla_a X^{\text{WF}}$ is a shorthand for $il_a X^{\text{WF}}(\mathbf{l})$.

The gradient of the log-determinant of the covariance is a mean-field term denoted by g^{MF} . Indeed, since the gradient of a log-likelihood always vanishes in the mean at the true parameter values, and since the covariance does not depend on the data, we can write $g^{\text{MF}} = \langle g^{\text{OD}} \rangle_\alpha$, where the average is taken over realizations of the CMB for a fixed deflection field α . This g^{MF} corresponds to a *delensing-induced* mean field, where the source of anisotropy is the lensing estimate itself [17].

Once converged, the lensing map ϕ^{MAP} is normalized by a Wiener filter, as shown in [6]. The fiducial value of the Wiener filter is

$$\mathcal{W}_L = \frac{C_L^{\phi\phi, \text{fid}}}{C_L^{\phi\phi, \text{fid}} + 1/R_L}, \quad (2.6)$$

where the response R_L is obtained iteratively from the partially delensed CMB spectra and the delensing efficiency, following the procedure of [6, 18]. However, since the fiducial Wiener filter is only an isotropic approximation of the response of ϕ^{MAP} to the true lensing signal, we always correct this normalization with Monte Carlo simulations.

B. Cross-only iterative estimator

Let us now assume we can split the observed CMB map into a set of n maps X_i^{dat} , observed at different times. For simplicity, we assume the same noise properties in each split, so that the optimal co-add map is simply the mean

$$X^{\text{dat}} = \frac{1}{n} \sum_{i=1}^n X_i^{\text{dat}}. \quad (2.7)$$

The quadratic term of the log-likelihood in Eq. (2.3) is then given by

$$X^{\text{dat}, \text{T}} \text{Cov}_\alpha^{-1} X^{\text{dat}} = \frac{1}{n^2} \sum_{i,j} X_i^{\text{dat}, \text{T}} \text{Cov}_\alpha^{-1} X_j^{\text{dat}}, \quad (2.8)$$

and the quadratic gradient becomes

$$g_a^{\text{OD}} = \frac{1}{n^2} \sum_{i,j} [\bar{X}_i]^x [\mathcal{D}_\alpha \nabla_a X_j^{\text{WF}}]_x. \quad (2.9)$$

Our goal now is to perform an iterative lensing reconstruction which neglects the information coming from the autocorrelations of the split maps. Our approach is to remove all the autocorrelation terms from the quadratic gradient in Eq. (2.9). This gives the following expression for a cross-only quadratic gradient:

$$g_a^{\text{OD}, \times} = \frac{1}{n(n-1)} \sum_{i \neq j} [\bar{X}_i]^x [\mathcal{D}_\alpha \nabla_a X_j^{\text{WF}}]_x, \quad (2.10)$$

where we also rescaled the normalization by $n/(n-1)$ such that its expectation value over realizations of the CMB fields is equivalent to the co-add gradient of Eq. (2.9). This cross-only gradient is similar to the cross-only QE of [12], adapted to the MAP formalism with the inclusion of the lensing and delensing operators in the filtering.

In practice, this cross-only gradient is the only modification we make to the standard MAP estimator. Thus, the filtering operations, as well as the Hessian matrix estimate (following the L-BFGS scheme, as in [4]) used in the iterations, are unchanged with respect to the standard (co-add) MAP estimator. We reconstruct the cross-only lensing field $\phi^{\text{MAP}, \times}$ by modifying the `delensalot`² pipeline to rely on the cross-only gradient (2.10) during the iterations, instead of the full co-add one (2.4).

Comparing the cross-only gradient in Eq. (2.10) to the full gradient of Eq. (2.9), we see that the information lost by removing the autocorrelations scales as $\mathcal{O}(1/n)$. This means that in the limit of a large number of splits, there should be no significant loss of information in performing the cross-only iterative reconstruction.

The costliest operation, which cannot be avoided, is to compute the n Wiener-filtered maps. In principle this Wiener-filtering step is the same for all maps, as the deflected covariance Cov_α is the same. However, the matrix inversion cannot be computed exactly due to its high dimensionality, and we use the conjugate-gradient inversion method to estimate the X_i^{WF} . This inversion method

¹Formally, the gradient is evaluated with respect to the two Cartesian components, α_1 and α_2 , of the deflection field in the flat sky.

²<https://github.com/NextGenCMB/delensalot>.

does not store the matrix inverse, only the product $\text{Cov}_\alpha^{-1} X_i^{\text{dat}}$, so we cannot repeat the same operation for all the splits, and we have to perform n conjugate-gradient inversions at each iteration.

C. Loss function

Let us remark that the cross-only gradient in Eq. (2.10) can be obtained from a loss function, which mimics the likelihood of the MAP estimator but neglects the autocorrelations of the split maps. This loss function is given by

$$f(\alpha, X_\times^{\text{dat}}) = -\frac{1}{2} X_\times^{\text{dat}T} K_\alpha^{-1} X_\times^{\text{dat}} - \frac{1}{2} \ln \det \text{Cov}_\alpha - \frac{1}{2} \text{Tr}[\text{Cov}_\alpha^{-1} N] - \frac{1}{2} \sum_L \frac{|\phi_L|^2}{C_L^{\phi\phi}}, \quad (2.11)$$

where we introduced the split data vector

$$X_\times^{\text{dat}} = (X_1^{\text{dat}}, X_2^{\text{dat}}, \dots, X_n^{\text{dat}})^T \quad (2.12)$$

and the block matrix K_α^{-1} , consisting of $n \times n$ blocks, which is zero on the block diagonal and Cov_α^{-1} elsewhere,

$$K_\alpha^{-1} = \frac{1}{n(n-1)} \begin{pmatrix} 0 & \text{Cov}_\alpha^{-1} & \dots & \text{Cov}_\alpha^{-1} \\ \text{Cov}_\alpha^{-1} & 0 & \dots & \text{Cov}_\alpha^{-1} \\ \vdots & \vdots & \ddots & \vdots \\ \text{Cov}_\alpha^{-1} & \text{Cov}_\alpha^{-1} & \dots & 0 \end{pmatrix}. \quad (2.13)$$

The gradient of the quadratic term of the loss function in Eq. (2.11) with respect to the deflection field gives back the cross-only gradient in Eq. (2.10). The block matrix K_α^{-1} mimics the covariance matrix of the standard likelihood, but is not positive-definite, although it is invertible. So it is not a covariance matrix and the loss function we introduced is not a likelihood anymore. This means that this loss function might not have an extremum. However, as validated in Sec. III with simulations, in the limit of signal-dominated maps, the quadratic form becomes in practice positive-definite.

Indeed, assuming $X_i^{\text{dat}} = X^{\text{dat}} + n_i$, the quadratic product between the i and j maps, with $i \neq j$, is

$$\begin{aligned} X_i^{\text{dat}T} \text{Cov}_\alpha^{-1} X_j^{\text{dat}} &= X^{\text{dat}T} \text{Cov}_\alpha^{-1} X^{\text{dat}} \\ &+ X^{\text{dat}T} \text{Cov}_\alpha^{-1} n_j + n_i^T \text{Cov}_\alpha^{-1} X^{\text{dat}} \\ &+ n_i^T \text{Cov}_\alpha^{-1} n_j. \end{aligned} \quad (2.14)$$

When the noise is negligible with respect to the data, we see that the quadratic form becomes in practice positive-definite, as only the part $X^{\text{dat}T} \text{Cov}_\alpha^{-1} X^{\text{dat}}$ contributes.

The mean-field part—i.e., the second and third terms in Eq. (2.11)—is defined such that the loss function is kept unbiased, in the sense that the gradient of the loss function (without the prior) vanishes in the mean at the true lensing potential (see the derivation in the Appendix). The gradient of the extra term $\text{Tr}[\text{Cov}_\alpha^{-1} N]$ corresponds to the mean field sourced by the delensing of the noise, as discussed in detail in Ref. [17]. This delensed noise mean field comes from anisotropies created in the noise map when delensing the CMB. As can be inferred from Eq. (2.11), the cross-only estimator does not contain the noise contribution to the delensing-induced mean field, which is in principle included in the standard MAP mean field given by the gradient of $\ln \det \text{Cov}_\alpha$. In practice we will neglect the delensing-induced mean field in the iterations. In [17] we showed that this mean-field term does not impact the lensing reconstruction cross-correlation coefficient. It appears only to impact the normalization of the estimated lensing field. So, provided the normalization is correctly estimated with Monte Carlo simulations, we can safely neglect the delensing-induced mean field.

D. Summary of the iterative process

We summarize below the iterative procedure of the cross-only MAP:

- (1) Get an estimate of the lensing deflection field α_N . For the first step we take the cross-only QE of [12].
- (2) Compute the n Wiener-filtered delensed split maps using a conjugate-gradient inversion.
- (3) Estimate the split gradient $\mathbf{g}^{\text{QD},\times}$ from Eq. (2.10).
- (4) Add the prior gradient and obtain the total gradient $\mathbf{g}_N$ (we neglect the delensed mean-field gradient).
- (5) The Hessian curvature matrix H_N (i.e., the second derivative of the log-posterior) is updated following the L-BFGS scheme, as in [4].
- (6) We obtain the next lensing deflection field as $\alpha_{N+1} = \alpha_N + \lambda H_N \mathbf{g}_N$, where we take $\lambda = 0.5$ to improve convergence.

We denote by $\phi^{\text{MAP},\times}$ the cross-only lensing potential obtained at convergence. Its effective normalization is determined using a set of Monte Carlo simulations, following the same procedure as for the standard MAP estimator.

Finally, we can estimate the mean field coming from mask, noise, or other anisotropic features not due to lensing, by averaging $\phi^{\text{MAP},\times}$ on a set of realistic Monte Carlo simulations, such as in [18]. This mean field is then subtracted from the estimated lensing field.

In our iterative delensing procedure, we use the deflection field estimated from split maps, α_N , to delens all the maps at the next iteration. This approach does not eliminate all noise contractions in the quadratic gradient estimator: higher-order contractions involving the same maps can remain.

To illustrate this schematically at the first iteration, consider two split maps, X_1 and X_2 . The cross-only QE for the lensing deflection can be written as in Eq. (2.10) (for $\alpha = 0$),³

$$\alpha_{(0)}^{12} \sim \frac{1}{2} (\bar{X}_1 \nabla X_2^{\text{WF}} + \bar{X}_2 \nabla X_1^{\text{WF}}). \quad (2.15)$$

The delensed CMB map at the first iteration is then approximately

$$X_1^{\text{del}} \simeq X_1 - \alpha_{(0)}^{12} \cdot \nabla X_1. \quad (2.16)$$

Assuming the MAP estimate behaves like a QE applied to delensed maps, the first iteration of the deflection field will contain contractions of the form

$$\alpha_{(1)}^{12} \sim (X_1 - \alpha_{(0)}^{12} \cdot \nabla X_1)(X_2 - \alpha_{(0)}^{12} \cdot \nabla X_2). \quad (2.17)$$

This expression contains residual contractions such as $X_1 X_1$ and $X_2 X_2$ through $\alpha_{(0)}^{12}$, implying that the noise is not completely cancelled in the iterative process. In this context, this corresponds to the noise contribution to the mean field. In principle, one could devise a more optimal scheme in which pairs of maps X_i and X_j are delensed using a deflection field α_N that is independent of both maps i and j . However, we find that our simpler method already suppresses most of the mean field noise contractions effectively, as demonstrated in Sec. III B.

E. Power spectrum from splits

The cross-only QE power spectrum, as introduced in [12], is defined as the sum of the lensing field estimated from all possible pairs of maps, avoiding any autocorrelation,

$$C_L^{\text{QE},\times} = \frac{1}{n(n-1)(n-2)(n-3)} \sum_{i \neq j \neq k \neq l} C_L(\hat{\phi}_{ij}^{\text{QE}}, \hat{\phi}_{kl}^{\text{QE}}), \quad (2.18)$$

where $\hat{\phi}_{ij}^{\text{QE}}$ means a QE lensing field estimated from the pair of maps i and j .

However, for the cross-only MAP we introduced in Sec. II, even if the gradient $g^{\text{QD},\times}$ in Eq. (2.10) can be decomposed into a sum of quadratic terms, the $\phi^{\text{MAP},\times}$ obtained at convergence cannot be separated into a sum of quadratic terms with pairs of CMB maps. Indeed, at each iteration we apply the lensing and delensing operators, based on the current estimate of $\phi^{\text{MAP},\times}$, so the final lensing map is a complex high-order combination of the split maps.

Assuming we can decompose the CMB maps into n splits, it is possible to estimate $\phi^{\text{MAP},\times}$ from any subset of these n splits. For example, we can take a subset a that contains k maps to estimate $\phi_a^{\text{MAP},\times}$, and the complementary subset b that contains the $n - k$ other maps to estimate $\phi_b^{\text{MAP},\times}$. The power spectrum $C_L(\phi_a^{\text{MAP},\times}, \phi_b^{\text{MAP},\times})$ will avoid repetitions in the pairs of maps.

It follows that, if we have n splits, we can estimate the lensing power spectrum with any combinations of the two subsets of maps of size k and $n - k$, for all k ,

$$C_L^{\text{MAP},\times} = \frac{1}{N_{\text{sets}}} \sum_{a,b} C_L(\phi_a^{\text{MAP},\times}, \phi_b^{\text{MAP},\times}), \quad (2.19)$$

where N_{sets} is the number of sets to average over. The number of subsets scales as $\mathcal{O}(2^n)$, which becomes quickly infeasible for large n . In a simple and more realistic scenario where $n = 4$, there are only six possible subsets, each involving a pair of maps. The cross-only MAP power spectrum is then simply

$$C_L^{\text{MAP},\times} = \frac{1}{3} \left[C_L(\phi_{12}^{\text{MAP},\times}, \phi_{34}^{\text{MAP},\times}) + C_L(\phi_{13}^{\text{MAP},\times}, \phi_{24}^{\text{MAP},\times}) + C_L(\phi_{14}^{\text{MAP},\times}, \phi_{23}^{\text{MAP},\times}) \right], \quad (2.20)$$

which is the same form as for the cross-only QE power spectrum.

For the standard quadratic estimator, the estimated lensing spectra contain the signal we want to measure plus biases due to other contractions in the four-point function of the CMB. Schematically, we have

$$C_L^{\hat{\phi}\hat{\phi}} = \hat{C}_L^{\phi\phi} + N_L^{(0)} + N_L^{(1)}, \quad (2.21)$$

with $N_L^{(0)}$ the disconnected four-point function and $N_L^{(1)}$ the connected four-point function that is first order in $C_L^{\phi\phi}$.

The power spectrum of the MAP also contains biases, noted $N_L^{(0),\text{MAP}}$ and $N_L^{(1),\text{MAP}}$ by analogy with the QE. In principle these are not only four-point functions of the CMB maps anymore, as evident from the structure of the first iteration in Eq. (2.17). However, we can predict these biases using partially delensed CMB spectra in the standard expression of the QE biases, as shown in [6].

For the cross-only MAP power spectrum, the disconnected contractions in the cross-power spectrum $C_L(\phi_{12}^{\text{MAP},\times}, \phi_{34}^{\text{MAP},\times})$ that contributes to the $N_L^{(0)}$ bias, such as $X_1 X_3$ and $X_2 X_4$, are, by construction, free of noise contributions, since they avoid any autocorrelations between the split maps. In Sec. III A we will confirm with noiseless simulations that the cross-only MAP $N_L^{(0),\times}$ bias is indeed free from noise contributions.

³There are implicit $\mathcal{D}_{\alpha=0}$ map-synthesis operators in the Wiener-filtered fields here, which we suppress to avoid clutter.




FIG. 1. Biases of the MAP and cross-only MAP estimators for our three experimental configurations. The mean biases obtained from 200 simulations with the MAP and cross-only MAP estimators are shown as the circles and squares, respectively. The solid and dotted lines are the predictions of the $N_L^{(0)}$ and $N_L^{(1)}$ biases for the MAP and the cross-only MAP estimators, respectively. The dashed lines are the residual biases obtained from 100 noiseless simulations with the standard (co-add) MAP estimator for each experimental configuration. The black solid line is the fiducial CMB lensing power spectrum. We show ACT-like temperature (orange), SO-like for temperature and for both temperature and polarization (purple and sky blue), and CMB-S4-like polarization (teal). We see that the CMB-S4 polarization configuration benefits the most from the bias reduction from the split estimator, while when the CMB fields are signal dominated, such as for the SO temperature case, the bias is less significantly reduced. Note that this bias reduction does not correspond to a reduction in the variance of the bandpower estimates.

We predict analytically the cross-only MAP $N_L^{(0),\times}$ bias by setting to zero the noise spectra in the standard $N_L^{(0),\text{MAP}}$ bias procedure. The $N_L^{(1)}$ bias is assumed to be independent of the noise in the CMB maps, as is the case for the QE. So we keep the same expression for the standard MAP and the cross-only MAP. The sums of these biases are shown in Fig. 1 as the solid and dotted lines for the standard and cross-only MAP estimators, respectively.

III. VALIDATION WITH SIMULATIONS

We now test our cross estimator and demonstrate that it converges to a close-to-optimal lensing field while being able to reduce both the lensing reconstruction noise at high L and the mean-field bias at low L .

A. Isotropic noise

We simulate CMB temperature and polarization full-sky maps, lensed by a Gaussian lensing potential field, and add Gaussian noise realizations. We consider three different experimental configurations, mimicking current and upcoming surveys, as listed in Table I. We make 200 simulations of the CMB for each experimental configuration. To mimic the split of the data into four maps, for each simulation we generate four Gaussian noise fields, with the noise levels increased by a factor of 2 in each map.

We reconstruct the lensing potential using the temperature maps for the ACT-like configuration, either the temperature only, or both temperature and polarization maps, for the

Simons Observatory (SO)-like configuration, and only the polarization maps for the CMB-S4-like configuration. For each of the 200 simulations, and each experiment, we reconstruct lensing with the QE and MAP estimator and with the cross-only QE and cross-only MAP estimator. For the MAP estimators, we run five iterations on each map, as we found that these achieve sufficient convergence. We estimate the effective normalization of ϕ^{MAP} and $\phi^{\text{MAP},\times}$ using the average of the cross-correlation of the 200 reconstructed maps with the true lensing fields, as in [6].

We show in Fig. 1 the biases on the lensing power spectrum for the MAP and cross-only MAP estimators, averaged over the 200 simulations. The circles and squares are the binned average of the lensing power spectra from ϕ^{MAP} and $\phi^{\text{MAP},\times}$, where we subtracted the input lensing spectrum. The predictions for the MAP biases, $N_L^{(0)} + N_L^{(1)}$,

TABLE I. Experimental configurations for the simulations. We consider configurations similar to ACT, SO, and CMB-S4. The polarization noise levels are given by $\Delta P = \sqrt{2}\Delta T$. The $\ell_{\min}$ and $\ell_{\max}$ are the CMB multipole range used for estimating the lensing potential field.

	Noise level ΔT [$\mu\text{K-arcmin}$]	Beam [arcmin]	$\ell_{\min}$	$\ell_{\max}$
ACT	10	3	40	3000
SO	6	1.4	100	3000
CMB-S4	1	1	40	4000




FIG. 2. Ratio of the CMB lensing power spectrum variance, for the cross-only estimators over the standard (co-add) MAP estimator, obtained from 200 full-sky simulations and binned in bandpowers. We show the cross-only QE (dashed lines) and cross-only MAP (solid lines), over the co-add MAP, for our three experimental configurations. For the SO-like case we show both the temperature only and the temperature and polarization estimators. The variance of the cross-only estimators are always higher than the co-add MAP estimator, as expected. We see that the increase in variance of the cross-only QE is greatly reduced for the cross-only MAP estimator, in particular for surveys including polarization. Comparing the cross-only MAP to the co-add MAP, we see that the increase in variance is below 5% for the ACT-like and SO-like configurations. For the CMB-S4 case, at multipoles $L < 1000$, where most of the signal-to-noise is accumulated, the degradation is still negligible.

are shown with solid lines, and the cross-only MAP biases, $N_L^{0,\times} + N_L^1$, are shown as the dotted lines.

The noise bias is reduced by a factor of around 5 for the CMB-S4 polarization case. The reduction is less significant in the SO temperature case. This is expected, as the CMB temperature is totally signal dominated in the range of multipoles we considered for the SO configuration (up to $\ell = 3000$), so the bias reduction from cancellation of the noise autocorrelations is negligible. We see that there is a good agreement between the predicted and estimated noise biases for both the MAP and cross-only MAP estimators. The prediction for cross-only MAP with the CMB-S4 configuration is around 10% lower than the simulation results for $L \sim 1500$, and around 50% lower for $L \sim 200$. In a realistic analysis, this small offset would be captured by a realization-dependent bias subtraction [RD- $N_L^{(0)}$] [19], and by a Monte Carlo correction applied on the lensing bandpowers, but neither are applied here.

To directly isolate the impact of noise, we generate 100 noiseless CMB simulations and reconstruct the lensing potential using the standard MAP estimator, applying the fiducial noise level in the filtering corresponding to each experimental setup. The residual biases from these reconstructions, plotted as the dashed lines in Fig. 1, match precisely the noise biases observed in our cross-only MAP




FIG. 3. Correlation matrices for the lensing power spectrum $C_L^{\phi\phi}$ in 25 bins between $L = 2$ and $L = 3000$. We show the standard MAP in the upper-left triangles, and the cross-only MAP in the lower-right triangles, for the four experimental configurations considered. We see that in the CMB-S4 polarization configuration, the cross-only covariance is visibly more diagonal than the standard MAP. Note that no realization-dependent bias correction is made, which would significantly reduce the off-diagonal correlations.

reconstructions across all configurations. This agreement provides compelling evidence that our cross-only MAP estimator entirely cancels the noise contribution to the lensing power spectrum biases.

Even if the biases are reduced with the cross-only MAP estimator, this does not correspond to a reduction in the variance of the lensing bandpowers over our set of simulations. We compare in Fig. 2 the bandpower errors for the cross-only estimators over the standard co-add MAP estimator. We see that the MAP cross-only estimator has a lower increase in variance than the cross-only QE, in particular for the CMB-S4 polarization and the SO minimum-variance cases. But even for the ACT-like temperature configuration, the increase is reduced by a factor of 2 compared to that for the cross-only QE. With the cross-only MAP in the temperature reconstruction, the increase in errors compared to the co-add MAP is negligible. For the minimum-variance SO-like configuration, the increase in variance of the cross-only MAP bandpowers compared to the co-add MAP ones is below 5% for all multipoles. In the CMB-S4-like polarization case, the errors for the cross-only MAP estimator are around 10% larger than the co-adds at high L , but below $L = 1000$ (which in practice achieves most of the signal-to-noise ratio) the increase is below 5% and thus negligible.

Figure 3 shows the power spectrum correlation matrices, i.e., the covariance matrices normalized by their diagonals,

for the MAP and cross-only MAP estimators. These are estimated from our 200 full-sky isotropic-noise simulations. Upper-left triangles are the standard MAP and the lower-right are the cross-only MAP estimators. We see that the correlations from the cross-only estimator are similar to those of the standard MAP estimator, and even slightly more diagonal for the CMB-S4 polarization configuration. By analogy with the QE, this could come from the reduction of the two-point function for some of the four products of two-point contractions that enter the dominant contribution to the off-diagonal covariance. The relevant terms are given in Eq. (46c) of [19]. The gain from not having noise in some of these is larger for polarization than temperature since the latter is almost signal-dominated.

Note that here we do not subtract realization-dependent forms of the bias corrections. In practice this would greatly reduce the nondiagonal correlations [6,19].

B. Anisotropic noise and mean field

Anisotropies in the CMB maps not due to lensing, such as anisotropic noise patterns or the mask of the survey, can bias the reconstructed lensing field. This contribution, called the mean field, is in general estimated from survey-specific Monte Carlo simulations, and subtracted from the measured lensing map. This requires the simulations to capture the survey strategy and atmospheric noise correlations accurately. In this section, we investigate the mean field due to time-varying noise (such as instrumental or atmospheric noise), which would create such anisotropies in the CMB maps. Since this noise is time-varying, the cross-only QE built from map splits with independent noise realizations should be able to cancel exactly the noise contribution to the mean field [12].

We generate maps with highly anisotropic noise variance based on the Planck simulations of noise in the 100 GHz channel. An example noise map is shown in Fig. 4. We rescale this map such that its variance is equal to the Gaussian noise level of the simulation settings we consider.




FIG. 4. One realization of the noise map used for the simulations, following the Planck scanning strategy, scaled such that its standard deviation is of 10 $\mu\text{K-arcmin}$.

For each CMB simulation, we generate four map splits by adding independent realizations of the anisotropic noise.

The mean field is estimated from a set of 100 simulations, where we reconstruct the QE, the MAP, and the cross-only MAP lensing fields. Since we want to test the nonideal case where we do not know the noise maps, we assume isotropic noise in the filtering operations of the QE and MAP estimators. To eliminate bias due to common simulations in each averaged mean-field estimate, the mean-field power spectrum is estimated from averaging two sets of 50 reconstructed maps, and taking the cross-correlation between these two sets of maps,

$$C_L^{\phi^{\text{MF}}} = C_L(\phi^{\text{MF1}}, \phi^{\text{MF2}}), \quad (3.1)$$

with ϕ^{MF1} the mean field obtained from averaging the reconstructed lensing field from the first set of simulations and ϕ^{MF2} from the second set.

For the cross-only estimator we consider four split maps, which gives six pairs of CMB maps to estimate the lensing field from. We estimate the lensing mean field by averaging the cross-only MAP reconstructed for each pair of maps. We estimate the mean-field power spectrum following the expression of Eq. (2.20) and computing the cross-spectra with a different set of simulations on each leg to avoid biases,

$$C_L^{\text{MF},\times} = \frac{1}{36} \sum_{i \neq j, k \neq l} C_L(\langle \phi_{ij}^{\text{MAP},\times} \rangle_a, \langle \phi_{kl}^{\text{MAP},\times} \rangle_b), \quad (3.2)$$

where the $\langle \rangle_a$ and $\langle \rangle_b$ are for averaging over a different subset, each with 50 simulations.

In Fig. 5 we show the mean-field power spectrum for the QE, the standard MAP, and the cross-only MAP. The cross-only QE mean field is zero by construction since we are considering full-sky observations. The right panel shows the mean field for the CMB-S4-like polarization configurations. The MAP mean field is already an order of magnitude lower than the QE mean field, and we see that cross-only MAP mean field is completely negligible.

In the left and right panels of Fig. 5 we show the mean field for the ACT-like and SO-like temperature configurations. The QE and MAP mean fields are very similar since in these configurations the iterative estimator is almost equivalent to the QE. However, we see that the cross-only MAP reduces the mean-field power by about an order of magnitude in both cases. This mean field is not totally nulled, contrary to the cross-only QE. This is likely due to the fact that during the iterative process, there are some nontrivial combinations of maps involved, as described schematically in Eq. (2.17).

We show as dotted lines the predicted mean-field power spectra arising from noise anisotropies for both the standard QE and MAP estimators. These predictions are obtained using the mean-field response derived in Appendix B of [9],




FIG. 5. Power spectra of the mean field for the QE (orange), standard MAP (blue), and cross-only MAP (teal), estimated from averaging over 100 simulations with anisotropic noise patterns. The QE cross-only mean field is zero by construction. For reference we show the fiducial lensing power spectrum (solid black line) and the $N_L^{(0),\text{MAP}}$ bias (dashed black line). The left panel is for the ACT-like configuration with temperature, the central panel is for a SO-like configuration with temperature, and the right panel is for CMB-S4 with polarization. We see that the cross-only MAP is able to reduce the mean-field term, in particular in the polarization CMB-S4 configuration. The mean field does not disappear in the temperature only reconstructions, but it is reduced by an order of magnitude for both ACT and SO. We also show the predicted mean-field power spectrum due to noise anisotropies as the dotted and dashed lines, for standard QE and MAP estimators, respectively.

generalized here to include the polarization-based estimator. For the MAP estimator, the mean-field response is computed using delensed CMB spectra (obtained via the iterative procedure employed in the calculation of the $N_L^{(0)}$ bias) in place of the lensed spectra. These predictions are found to be in good agreement with the mean-field power spectra estimated from simulations.

IV. CONCLUSION

In this work we modified the standard maximum *a posteriori* lensing estimator to allow for cross-only lensing reconstruction from CMB maps with independent noise. This allows removal of most of the contributions from the autocorrelations of the noise in the CMB maps. We showed that, despite the lack of formal guarantee of convergence, it is possible to maximize the loss function we introduced and reconstruct a lensing field that is close to optimal. This is especially true when the signal dominates the noise, as in our SO-like temperature only configuration.

As for the cross-only QE, our cross-only MAP estimator is able to reduce the noise biases, as well as the mean-field bias coming from noise. The reduction of these biases is most important for polarization-based lensing reconstruction, such as our CMB-S4-like polarization only

configuration, since the noise is large on the small scales of interest for the lensing. The cross-only estimator also reduces the off-diagonal correlations between the CMB lensing bandpowers.

We did not investigate the implementation of realization-dependent bias subtraction [RD- $N_L^{(0)}$, see [9] and references therein]. This method provides a first-order accurate estimate of the $N_L^{(0)}$ bias in the presence of unmodeled survey anisotropies, and reduces correlations between multipoles in the debiased CMB lensing spectra. In principle, one could start with the formalism of the cross-only QE RD- $N_L^{(0)}$ introduced in [12] and extend it to the RD- $N_L^{(0),\text{MAP}}$ framework of [6,18], thereby deriving a realization-dependent debiaser for our cross-only MAP estimator. However, we leave this extension for future work. Importantly, we have shown that the covariance of the cross-only MAP is slightly more diagonal than that of the standard MAP, and the cross-only $N_L^{(0)}$ bias is lower. This suggests that the impact of realization-dependent bias corrections may be less significant in our case.

The cross-only estimator we introduced is close to the optimal lensing reconstruction and robust to the complex noise properties that arise from ground-based observations. It will be a useful estimator to validate a standard MAP lensing analysis. Indeed, comparing the reconstruction between different splits of the data, such as in [9], is important to check for potential unknown systematics.

Finally, our formalism opens the door to new extensions of the QE tools into the MAP formalism. For instance, following the gradient-leg cleaning estimator from [20,21], we should be able to define an iterative estimator where one leg (the one with the Wiener filtering) of the quadratic gradient has been foreground cleaned, while the other leg retains the full information. One could also extend the halo lensing estimator of [22], with one leg with a lower ℓ_{max} filtering.

Our work suggests that the MAP formalism is robust and sufficiently flexible that the tools developed for the QE can be extended to the MAP (see also [23,24]). This is a promising result for the next generation of CMB surveys, where the MAP formalism will be necessary to extract the maximum information from the data.

ACKNOWLEDGMENTS

The authors thank Niall MacCrann and Jack Kwok for helpful discussions during the early stages of this work. L. L. acknowledges support from the Swiss National Science Foundation through the Postdoc.Mobility fellowship (Grant No. P500PT_21795). B. D. S. acknowledges support from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant Agreement No. 851274). A. C. acknowledges support from the STFC (Grant Numbers ST/W000977/1 and ST/X006387/1). J. C. acknowledges

support from a SNSF Eccellenza Professorial Fellowship (No. 186879). G. S. F. is supported by Lawrence Berkeley National Laboratory and the Director, Office of Science, Office of High Energy Physics of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. Part of this work uses resources provided by the Cambridge Service for Data Driven Discovery (CSD3) operated by the University of Cambridge Research Computing Service [25], provided by Dell EMC and Intel using Tier-2 funding from the Engineering and Physical Sciences Research Council (capital grant EP/T022159/1), and DiRAC funding from the Science and Technology Facilities Council [26].

DATA AVAILABILITY

The data are not publicly available. The data are available from the authors upon reasonable request.

APPENDIX: MEAN FIELD OF THE CROSS-ONLY ESTIMATOR

We derive here the expression for the additional mean-field term, $-\text{Tr}[\text{Cov}_\alpha^{-1}N]/2$, that we introduce in the loss function (2.11) for our split-only estimator. Assuming we split the data into n split maps, we rewrite the block matrix K_α^{-1} from Eq. (2.13) as a Kronecker product

$$K_\alpha^{-1} = \frac{1}{n(n-1)} (J_n - I_n) \otimes \text{Cov}_\alpha^{-1}. \quad (\text{A1})$$

$$\begin{aligned} \left\langle -\frac{1}{2} X_\times^{\text{dat}T} K_\times^{-1} X_\times^{\text{dat}} \right\rangle_{\tilde{\alpha}} &= -\frac{1}{2} \text{Tr}[K_\times^{-1} \widetilde{\text{Cov}}(X_\times^{\text{dat}})] \\ &= -\frac{1}{2} \text{Tr} \left[\left(\frac{1}{n(n-1)} (J_n - I_n) \otimes \text{Cov}_\alpha^{-1} \right) (J_n \otimes (\widetilde{\text{Cov}}_\alpha - N) + I_n \otimes N_\times) \right] \\ &= -\frac{1}{2} \frac{1}{n(n-1)} \text{Tr}[(J_n - I_n) J_n \otimes \text{Cov}_\alpha^{-1} (\widetilde{\text{Cov}}_\alpha - N) + (J_n - I_n) \otimes \text{Cov}_\alpha^{-1} N_\times] \\ &= -\frac{1}{2} \frac{1}{n(n-1)} \text{Tr}[(n-1) J_n \otimes (\text{Cov}_\alpha^{-1} \widetilde{\text{Cov}}_\alpha - \text{Cov}_\alpha^{-1} N)] \\ &= -\frac{1}{2} \text{Tr}[\text{Cov}_\alpha^{-1} \widetilde{\text{Cov}}_\alpha] + \frac{1}{2} \text{Tr}[\text{Cov}_\alpha^{-1} N]. \end{aligned} \quad (\text{A4})$$

Taking the gradient of this expectation with respect to α , and then setting $\tilde{\alpha} = \alpha$, is equivalent to taking the gradient of $-X_\times^{\text{dat}T} K_\times^{-1} X_\times^{\text{dat}}/2$ first and then averaging over data lensed by fixed α . Applied to the right of Eq. (A4), this gives

$$\begin{aligned} &-\frac{1}{2} \text{Tr}[(\nabla_\alpha \text{Cov}_\alpha^{-1}) \text{Cov}_\alpha] + \frac{1}{2} \text{Tr}[(\nabla_\alpha \text{Cov}_\alpha^{-1}) N] \\ &= \frac{1}{2} \nabla_\alpha (\ln \det \text{Cov}_\alpha) + \frac{1}{2} \nabla_\alpha \text{Tr}[\text{Cov}_\alpha^{-1} N]. \end{aligned} \quad (\text{A5})$$

Here, J_n is the square matrix of size n with all elements equal to unity and I_n is the identity matrix of size n . The covariance between two split maps X_i and X_j is

$$\begin{aligned} \text{Cov}(X_i, X_j) &= \mathcal{B} \mathcal{D}_\alpha \mathcal{C}^{\text{unl}} \mathcal{D}_\alpha^\dagger \mathcal{B}^\dagger + \delta_{ij} N_\times \\ &= \text{Cov}_\alpha - N + \delta_{ij} N_\times, \end{aligned} \quad (\text{A2})$$

with $N_\times$ the noise covariance of the split maps and δ_{ij} the Kronecker delta. The main assumption is that the noise is uncorrelated between splits. For simplicity we also assume that the splits receive an equal noise contribution given by $N_\times$, so that $N_\times = nN$, but that assumption can be relaxed. The covariance of the split data vector $X_\times^{\text{dat}}$ from Eq. (2.12) is then

$$\text{Cov}(X_\times^{\text{dat}}) = J_n \otimes (\text{Cov}_\alpha - N) + I_n \otimes N_\times. \quad (\text{A3})$$

We want the gradient of the loss function (minus the prior term) to vanish at any α when averaged over data lensed by that same α . More generally, let us distinguish between a point in parameter space, α , and the actual lensing of the data, which we denote with a tilde. Averaging the quadratic part of the loss function defined in Eq. (2.11) over CMB realizations with fixed lensing field $\tilde{\alpha}$ we get

These contributions to the mean gradient are compensated by the second and third terms, respectively, in the loss function (2.11).

The first term on the right of Eq. (A5) is nothing other than the mean field of the standard MAP estimator. It corresponds to the mean field induced by the delensing, and it is discussed in detail in [17]. The other term, $\nabla_\alpha \text{Tr}[\text{Cov}_\alpha^{-1}N]$, corresponds to subtracting from the delensing-induced mean field the part that comes directly from the delensing of the noise. Thus, our cross-only estimator is, in principle, insensitive to the mean field created by the delensing of the noise.

- [1] Antony Lewis and Anthony Challinor, Weak gravitational lensing of the CMB, *Phys. Rep.* **429**, 1 (2006).
- [2] Takemi Okamoto and Wayne Hu, CMB lensing reconstruction on the full sky, *Phys. Rev. D* **67**, 083002 (2003).
- [3] Christopher M. Hirata and Uros Seljak, Reconstruction of lensing from the cosmic microwave background polarization, *Phys. Rev. D* **68**, 083002 (2003).
- [4] Julien Carron and Antony Lewis, Maximum *a posteriori* CMB lensing reconstruction, *Phys. Rev. D* **96**, 063510 (2017).
- [5] Marius Millea and Uros Seljak, Marginal unbiased score expansion and application to CMB lensing, *Phys. Rev. D* **105**, 103531 (2022).
- [6] Louis Legrand and Julien Carron, Lensing power spectrum of the cosmic microwave background with deep polarization experiments, *Phys. Rev. D* **105**, 123519 (2022).
- [7] D. M. Regan, E. P. S. Shellard, and J. R. Fergusson, General CMB and primordial trispectrum estimation, *Phys. Rev. D* **82**, 023520 (2010).
- [8] Toshiya Namikawa, Duncan Hanson, and Ryuichi Takahashi, Bias-hardened CMB lensing, *Mon. Not. R. Astron. Soc.* **431**, 609 (2013).
- [9] N. Aghanim *et al.*, Planck 2018 results. VIII. Gravitational lensing, *Astron. Astrophys.* **641**, A8 (2020).
- [10] Julien Carron, Mark Mirmelstein, and Antony Lewis, CMB lensing from Planck PR4 maps, *J. Cosmol. Astropart. Phys.* **09** (2022) 039.
- [11] Frank J. Qu *et al.* (ACT Collaboration), The Atacama Cosmology Telescope: A measurement of the DR6 CMB lensing power spectrum and its implications for structure growth, *Astrophys. J.* **962**, 112 (2024).
- [12] Mathew S. Madhavacheril, Kendrick M. Smith, Blake D. Sherwin, and Sigurd Naess, CMB lensing power spectrum estimation without instrument noise bias, *J. Cosmol. Astropart. Phys.* **05** (2021) 028.
- [13] Blake D. Sherwin and Sudeep Das, CMB lensing—power without bias, [arXiv:1011.4510](https://arxiv.org/abs/1011.4510).
- [14] Delon Shen, Emmanuel Schaan, and Simone Ferraro, CMB lensing power spectrum without noise bias, *Phys. Rev. D* **110**, 043523 (2024).
- [15] Sebastian Belkner, Julien Carron, Louis Legrand, Caterina Umiltà, Clem Pryke, and Colin Bischoff (CMB-S4 Collaboration), CMB-S4: Iterative internal delensing and r constraints, *Astrophys. J.* **964**, 148 (2024).
- [16] Martin Reinecke, Sebastian Belkner, and Julien Carron, Improved cosmic microwave background (de-)lensing using general spherical harmonic transforms, *Astron. Astrophys.* **678**, A165 (2023).
- [17] Louis Legrand and Julien Carron, Characterizing CMB noise anisotropies from CMB delensing, *Phys. Rev. D* **112**, 083538 (2025).
- [18] Louis Legrand and Julien Carron, Robust and efficient CMB lensing power spectrum from polarization surveys, *Phys. Rev. D* **108**, 103516 (2023).
- [19] Duncan Hanson, Anthony Challinor, George Efstathiou, and Pawel Bielewicz, CMB temperature lensing power reconstruction, *Phys. Rev. D* **83**, 043005 (2011).
- [20] Mathew S. Madhavacheril and J. Colin Hill, Mitigating foreground biases in CMB lensing reconstruction using cleaned gradients, *Phys. Rev. D* **98**, 023534 (2018).
- [21] Omar Darwish *et al.*, The Atacama Cosmology Telescope: A CMB lensing mass map over 2100 square degrees of sky and its cross-correlation with BOSS-CMASS galaxies, *Mon. Not. R. Astron. Soc.* **500**, 2250 (2020).
- [22] Wayne Hu, Simon DeDeo, and Chris Vale, Cluster mass estimators from CMB temperature and polarization lensing, *New J. Phys.* **9**, 441 (2007).
- [23] Julien Carron, On likelihood-based analysis of the gravitationally (de)lensed CMB, [arXiv:2502.02399](https://arxiv.org/abs/2502.02399).
- [24] Omar Darwish, Optimal joint reconstruction from CMB observations: Application to cosmic birefringence, patchy reionization and CMB lensing, *Phys. Rev. D* **112**, 023514 (2025).
- [25] <http://www.csd3.cam.ac.uk>.
- [26] <http://www.dirac.ac.uk>.