

Characterizing CMB noise anisotropies from CMB delensing

Louis Legrand^{1,2} and Julien Carron³

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



(Received 12 May 2025; accepted 19 September 2025; published 21 October 2025)

Undoing the effect of gravitational lensing on the cosmic microwave background (CMB) (“delensing”) is essential in shaping constraints on weak signals limited by lensing effects on the CMB, for example on a background of primordial gravitational waves. Removing these anisotropies induced by large-scale structures from the CMB maps also generally helps our view of the primordial Universe by sharpening the acoustic peaks and the damping tail. However, practical implementations of delensing transfer parts of these anisotropies to the noise maps. This will induce a new large scale “mean-field” bias to any anisotropy estimator applied to the delensed CMB, and this bias directly traces large-scale structures. This paper analytically quantifies this delensed noise mean field and its impact on quadratic and likelihood-based lensing estimators. We show that for Simons-Observatory-like surveys, this mean-field bias can reach 15% in cross-correlation with large-scale structures if unaccounted for. We further demonstrate that this delensed noise mean field can be safely neglected in likelihood-based estimators without compromising the quality of lensing reconstruction or *B*-mode delensing, provided the resulting lensing map is properly renormalized.

DOI: [10.1103/dcj9-hhl4](https://doi.org/10.1103/dcj9-hhl4)

I. INTRODUCTION

The light of the cosmic microwave background (CMB) is deflected by the gravitational lensing due to the large scale structures [1]. Lensing thus deforms and blurs the image of the primordial CMB. This is of particular importance when trying to observe the primordial *B* modes, sourced by the gravitational waves created by inflation: constraints on the tensor to scalar ratio are limited by the variance from the secondary *B* modes created by lensing [2,3]. To improve constraints on the tensor scalar ratio, delensing of the CMB maps is key [4–7]. Delensing can also help in sharpening the peaks of the CMB temperature and polarization *E* modes spectra, which can improve the constraints on cosmological parameters [8,9]. Finally, delensing is at the core of the more efficient, likelihood-based CMB lensing methodology including maximum *a posteriori* (MAP) reconstruction [5,10–13] or marginal unbiased score expansion (MUSE) [14–16].

By construction, delensing will reduce the statistical anisotropies in the observed CMB that were created by lensing. However, it will generate anisotropies in the CMB noise maps.

Indeed, even when assuming Gaussian and statistically isotropic noise maps, the delensing will deform this noise map in such a way that it will become statistically anisotropic.

These features in the noise maps can be characterized in CMB lensing terminology as a *mean field*—a source of anisotropy potentially contaminating the signal of primary interest. Since this mean field is introduced by delensing the noise maps, we will call it a *delensed noise mean field*.

The scope of this article is to characterize this delensed noise mean field. In Sec. II we start by investigating qualitatively the delensed noise mean field and obtain a perturbative expression for it. Then in Sec. III we investigate the impact of the delensed noise mean field on the temperature only MAP estimator, and compare different ways of taking into account its contribution. We then check the impact of the delensed noise mean field in delensing the *B* modes in Sec. III E, and we conclude in Sec. IV. In the Appendix we detail the derivation of the delensed noise mean-field response.

II. DELENSED NOISE MEAN FIELD

In this Section we provide a useful if slightly qualitative path toward the main results, neglecting all the complications of real data. This is justified later, where we demonstrate rigorously that they also hold in more general contexts, including that of likelihood-based, “beyond-quadratic-estimator” lensing estimation in the presence of any sort of nonidealities.

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.

A useful manner to think about delensing is simply remapping of the coordinate: let $X^{\text{dat}}(\mathbf{x}) = X^{\text{unl}}(\mathbf{x} + \boldsymbol{\alpha}(\mathbf{x})) + n(\mathbf{x})$ be the observed CMB Stokes fields plus noise, where $\boldsymbol{\alpha}$ is the lensing deflection field. We assume in this section for simplicity that the maps were successfully beam deconvolved. Provided with any tracer of the lensing field $\hat{\boldsymbol{\alpha}}$, one can simply attempt to remap the points

$$X^{\text{del}}(\mathbf{x}) \equiv X^{\text{dat}}(\mathbf{x} - \hat{\boldsymbol{\alpha}}(\mathbf{x})) \quad (2.1)$$

(possibly after filtering the data, the tracer, or both). This is for example what the first delensing attempts by [4, 17] did.¹ Reconstruction of the unlensed CMB based on a likelihood model will perform the delensing in a more subtle manner, but we will see that the end result will be equivalent for our purposes—we stick to this remapping picture in this Section.

Under this operation, the noise component is obviously also transformed, according to

$$n(\mathbf{x}) \rightarrow n(\mathbf{x} - \hat{\boldsymbol{\alpha}}(\mathbf{x})), \quad (2.2)$$

and anisotropies will be induced in its covariance. Regions that were magnified by large scale lenses will be compressed to a smaller area, while preserving the surface brightness. Assuming the original instrumental noise was uncorrelated between pixels, this will reduce by some amount the effective local noise. The opposite happens in regions that were demagnified, where delensing will make the local noise go a bit up, by inflating the area by some amount.

We illustrate this in Fig. 1, where we show a patch of the effective temperature depth map from *Planck* [19,20], and the resulting delensed noise depth map. We see that the noise variance map is modulated by the convergence field (which is related to the deflection field, $\kappa = -\nabla \cdot \boldsymbol{\alpha}/2$), creating anisotropies in the delensed noise map that are tracing the lensing convergence.

Consider then a standard quadratic estimator (QE) from temperature [21,22], applied to these delensed maps

$$\begin{aligned} \bar{x}_{LM} \equiv & \frac{1}{2} (-1)^M \sum_{\ell_1 m_1 \ell_2 m_2} \begin{pmatrix} L & \ell_1 & \ell_2 \\ -M & m_1 & m_2 \end{pmatrix} \\ & \times W_{L\ell_1\ell_2}^x \bar{T}_{\ell_1 m_1}^{\text{del}} \bar{T}_{\ell_2 m_2}^{\text{del}} \end{aligned} \quad (2.3)$$

targeting some source of anisotropy x from the signal part of these maps, for example the residual lensing signal, or something else such as patchy reionization and cosmic birefringence [23–26]. Here the map $\bar{T}^{\text{del}}$ is the inverse-

¹Note that this naive approach can lead to suboptimal results, see [18] for a detailed discussion on delensing templates.




FIG. 1. Top panel shows a randomly selected 7 degrees by 7 degrees patch of *Planck* effective temperature depth map after spectral matching independent component analysis component separation [19,20]. Delensing using a Wiener-filtered lensing tracer with fidelity similar to that expected from the *Planck* reconstruction will result in an effective depth given by the middle panel, where the original stripes caused by the scan strategy are modulated by the deflection field magnification. The resulting fluctuations trace the convergence map of the deflection field shown in the bottom panel.

variance filtered delensed T^{del} map. Ignoring any non-idealities other than lensing in this Section, there are two main options. Either one may proceed in analogy to standard QE by ignoring the anisotropy in the covariance and writing

$$\bar{T}_{\ell m}^{\text{del}} = \frac{T_{\ell m}^{\text{del}}}{C_{\ell}^{\text{del}} + N_{\ell}}, \quad (2.4)$$

where C_{ℓ}^{del} is the delensed CMB signal spectrum and N_{ℓ} some choice of anisotropic noise spectrum. Or, one may take into account the fact that the CMB was delensed, and use instead “optimal-filtering,”

$$\bar{T}^{\text{del}} = (C^{\text{del}} + N_{\hat{\alpha}})^{-1} T^{\text{del}} \quad (2.5)$$

where $N_{\hat{\alpha}}$ is now a dense matrix, accounting for the delensing-induced anisotropies in the noise.

In both cases, these estimators $\bar{x}$ will respond to the noise anisotropy induced by $\hat{\alpha}$ in the delensed maps, besides the sought-after signal. As usual, we can decompose the deflection field into a gradient ($\kappa = -\Delta\phi/2$) and curl ($\omega = -\Delta\Omega/2$) component. We can then write schematically

$$\bar{x}_{LM} \ni \mathcal{R}_L^{\kappa, \text{MF}} \hat{\kappa}_{LM} + \mathcal{R}_L^{\omega, \text{MF}} \hat{\omega}_{LM} + \dots \quad (2.6)$$

where the response function are defined as $\mathcal{R}_L^{\kappa, \text{MF}} = \partial \bar{x}_{LM} / \partial \hat{\kappa}_{LM}$ and $\mathcal{R}_L^{\omega, \text{MF}} = \partial \bar{x}_{LM} / \partial \hat{\omega}_{LM}$. We have added the superscript MF to these responses since in quadratic estimator language they are “mean fields,” i.e., a bias in the estimated map $\bar{x}$, rather than sought-after anisotropic signals in the CMB maps.

The responses for the filtering choices (2.4) and (2.5) will be different. However, to leading order in $\hat{\alpha}$, one has simply

$$\begin{aligned} \mathcal{R}_L^{\kappa, \text{MF}}[\text{isotropic filtering, (2.4)}] \\ = -\mathcal{R}_L^{\kappa, \text{MF}}[\text{anisotropic filtering, (2.5)}] \end{aligned} \quad (2.7)$$

and similarly for ω . This is shown in details in the Appendix, but it is not too difficult to see why that can be: regions where the local noise variance is made higher (lower) by delensing will be accordingly down-weighted (up-weighted) by the optimal filtering (2.5), resulting eventually in an anisotropy of the opposite sign in the filtered map. From now on, the sign of our mean-field responses refer to optimal filtering (2.5). The reason is that this corresponds to the case of beyond-QE, likelihood-based lensing reconstruction, which always filters maps using all information available.

The response of a standard QE defined as in (2.3), but applied on temperature maps without delensing, is given by [21,22]

$$\mathcal{R}_L^{\kappa} = \frac{1}{2(2L+1)} \sum_{\ell_1 \ell_2} f_{L\ell_1 \ell_2}^{TT, \kappa} W_{L\ell_1 \ell_2}^x F_{\ell_1}^T F_{\ell_2}^T, \quad (2.8)$$

where $f^{TT, \kappa}$ is the response of the pair of fields TT to the lensing convergence, as given in Table 1 of [21], and F_{ℓ}^T is the inverse variance filtering approximated as diagonal in ℓ .

To obtain the delensed noise mean-field response $\mathcal{R}_L^{\kappa, \text{MF}}$, we make two modifications to this fiducial response (see the Appendix for a detailed derivation). First, we replace the filters by $F_{\ell} = 1/(C_{\ell}^{\text{del}} + N_{\ell})$, to match the delensing filtering as in Eqs. (2.4) and (2.5). Second, since this mean field is sourced by the noise, we replace the lensed CMB spectra (or the lensed gradient spectra, see [27,28]) entering the response f by the CMB noise spectra N_{ℓ} . The QE weights W depend on the choice of QE made. In Sec. III, with the MAP estimator, we will use delensed spectra C_{ℓ}^{del} in the weights W instead of the lensed (or lensed gradient) spectra.

Since we expect the delensed noise mean field to follow the deflection field, whose power spectrum peaks at large scales, $L \sim 30$, we now discuss the squeezed limit of the response. Let $\mathcal{R}_L^{\kappa, \text{MF}}$ and $\mathcal{R}_L^{\omega, \text{MF}}$ be the linear response of the convergence κ estimator mean field to κ , and of the field rotation ω estimator mean field to ω , respectively (as always, from parity arguments cross terms vanish). Since the only difference to a standard QE response is that the anisotropy source enters the noise, we can use as starting point the squeezed limit of the standard estimator and make the corresponding modifications. The standard estimator response at low L is [29,30]

$$\begin{aligned} \mathcal{R}_L^{\kappa} \sim \frac{1}{2} \sum_{\ell} \left(\frac{2\ell+1}{4\pi} \right) \mathcal{W}_{\ell}^2 \\ \times \left[\left(\frac{d \ln \ell^2 C_{\ell}}{d \ln \ell} \right)^2 + \frac{1}{2} \Gamma_L^2 \left(\frac{d \ln C_{\ell}}{d \ln \ell} \right)^2 \right] \end{aligned} \quad (2.9)$$

where $\mathcal{W}$ is the isotropic Wiener filter

$$\mathcal{W}_{\ell} \equiv \frac{C_{\ell}}{C_{\ell} + N_{\ell}}, \quad (2.10)$$

and

$$\Gamma_L^2 \equiv \frac{(L+2)(L-1)}{L(L+1)} \quad (2.11)$$

is the shear factor.

The first term in brackets is the information from local magnification (vanishing for scale-invariant spectra), and the second from shear (vanishing for white spectra). The L -dependent shear prefactor is only relevant at the very lowest multipoles. The mean-field response is found by realizing that in the general case, one factor of C_{ℓ} in the numerators is the fiducial spectrum used as QE weight, and the second

is the spectrum affected by the anisotropy. Hence, the correct result is then found by replacing one of the two C_ℓ by N_ℓ . Doing so, one obtains

$$\mathcal{R}_L^{KK, \text{MF}} \sim \frac{1}{2} \sum_\ell \left(\frac{2\ell+1}{4\pi} \right) \mathcal{W}_\ell (1 - \mathcal{W}_\ell) \times \left[\frac{d \ln \ell^2 C_\ell^{\text{del}}}{d \ln \ell} \frac{d \ln \ell^2 N_\ell}{d \ln \ell} + \frac{1}{2} \Gamma_L^2 \frac{d \ln C_\ell^{\text{del}}}{d \ln \ell} \frac{d \ln N_\ell}{d \ln \ell} \right] \quad (2.12)$$

with

$$\mathcal{W}_\ell \equiv \frac{C_\ell^{\text{del}}}{C_\ell^{\text{del}} + N_\ell}. \quad (2.13)$$

The factor $\mathcal{W}_\ell (1 - \mathcal{W}_\ell)$ in Eq. (2.12) selects now the CMB modes where noise is relevant, but not too noisy so as to be useless. The factor $(1 - \mathcal{W}_\ell)$ appears by replacing the C_ℓ by N_ℓ in the numerator of $\mathcal{W}_\ell$, and corresponds a Wiener filter of the noise.

Further, from the same arguments, the lensing rotation response from large lenses is identical to the shear-part of the lensing convergence response from large lenses,

$$\mathcal{R}_L^{\omega\omega, \text{MF}} \sim \frac{1}{2} \sum_\ell \frac{2\ell+1}{4\pi} \mathcal{W}_\ell (1 - \mathcal{W}_\ell) \times \left[\frac{1}{2} \Gamma_L^2 \frac{d \ln C_\ell^{\text{del}}}{d \ln \ell} \frac{d \ln N_\ell}{d \ln \ell} \right]. \quad (2.14)$$

The reason for this is as follows: in contrast to the (de)magnifying effect of a constant convergence, a constant rotation on a patch of the CMB temperature has no locally observable effect at all. Hence, it is the contribution of ω to the shear B mode responsible for Eq. (2.14), which is the same as that of κ to the shear E -mode contribution to Eq. (2.12).

For high-resolution experiments, a white noise model is often accurate over a wide range of scales. Hence, the limit of perfectly white noise is both simple and of some relevance. In this case only the convergence piece contributes, and can be simplified further² to the following result, with ℓ_{max} the maximum multipole used in the CMB temperature maps.

$$\mathcal{R}_L^{KK, \text{MF}} \sim -2 \sum_\ell \frac{2\ell+1}{4\pi} (\mathcal{W}_\ell^2 - \mathcal{W}_{\ell=\ell_{\text{max}}}), \quad (2.15)$$

²To get (2.15), we used $(1 - \mathcal{W}) d \ln C = d \ln \mathcal{W}$, which holds for white noise, and is integrated by part. This limit does match the estimate obtained independently combining B.9 and B.12 of [5], and improves slightly the one given in the main text of that reference.

$$\mathcal{R}_L^{\omega\omega, \text{MF}} = 0 \quad (\text{low } L\text{'s, white noise}). \quad (2.16)$$

This shows that as for the standard lensing response, the dependence on ℓ_{max} can be very strong. If the range of CMB multipoles used is large enough and enters the noise dominated regime, the response will saturate to $-2 \sum \mathcal{W}_\ell^2 (2\ell+1)/4\pi$. On the other hand, it is always zero in the regime where all the modes considered are noise-free ($\mathcal{W}_\ell \sim 1$), which is quite relevant for a deep experiment like CMB-S4 for example.

Realistic noise is of course not perfectly white, but grows sharply on scales below the resolution of the experiment. This can make the shear-type contribution very substantial. For calculations we use a standard Gaussian model for which

$$\ln N_\ell = \left(\frac{\sigma_{\text{FWHM}}}{8 \ln 2} \right)^2 \ell(\ell+1) + \text{const}, \quad (2.17)$$

where σ_{FWHM} is the beam full width at half maximum. The contribution per mode to the delensed noise mean-field response is shown in Fig. 2 for a Simons observatory like configuration.

Let us now discuss briefly the case of polarization. For lensing on intermediate scales, e.g., $\ell_{\text{max}} = 3000$, and low instrumental noise, it is well known that the EB estimator carries most of the signal. Large lenses create locally a EB spectrum proportional to $C_\ell^{EE} - C_\ell^{BB}$, see, e.g., [28]. Since




FIG. 2. Blue line shows the contribution per multipole to the large scale (here $L = 50$) delensed noise mean-field response to the lensing convergence map, for a Simons observatory like configuration, with a Gaussian beam of three arcmin and white noise. The squeezed approximation in (2.12) is a perfect match, with the orange curve highlighting the contribution from converging rather than shearing effects on the delensed noise map. The green line shows for comparison the standard lensing signal contributions to the TT lensing quadratic estimator (reduced by a factor 5). The delensed noise mean field is sourced from significantly smaller scales, where noise is relevant but not completely dominating either. This is for temperature reconstruction.




FIG. 3. Delensed-noise mean-field response for the polarized lensing quadratic estimator, for a configuration approaching crudely the planned CMB-S4 wide survey, with equal power in the E and B noise spectra. All curves were computed using the lensed CMB spectra. In contrast to standard lensing, the EB mean-field response is essentially zero (green), owing to the symmetry in unlensed E and B noise power, and EE (blue) dominate the response, followed by BB in orange. Dashed green shows the standard lensing response in the same configuration for comparison.

the primordial C_ℓ^{BB} spectrum is very weak, this makes this signal very significant, provided a sufficient number of modes is resolved. On the other hand, the EE estimator behaves similarly to temperature (the magnification and shearing effects discussed above), while the BB lensing estimator is near irrelevant.

Lensing acting on the noise maps is different, since the strong E - B asymmetry of the primordial CMB signal is lost. Since N^{EE} is comparable to N^{BB} , there is no creation of small-scale EB power, and the EB mean-field response will vanish on large scales. Both EE and BB will behave now like the temperature case (2.12). The relative weight of the mean-field response compared to the EB response is much weaker, so that the overall effect in low-noise polarization reconstruction is always very small. Likewise, this should only have a very small impact on the TE and TB estimators, which probes the unlensed TE spectrum. Figure 3 shows the polarization mean-field responses for a CMB-S4-like configuration.

III. IMPACT ON OPTIMAL-LENSING RECONSTRUCTION

We now investigate the role and impact of the delensed noise mean field on likelihood-based lensing reconstruction. First proposed by [10], usage of the CMB-lensing likelihood allows better (in principle, optimal) reconstruction of the CMB lensing deflections.

We start in Sec. III A discussing how this mean field enters the original approach of [10], which was revisited more recently by [11]. See [5] for a detailed description of the resulting algorithm in spherical geometry.

Reference [14] introduced the alternative approach of using the joint primordial CMB and CMB lensing likelihood, whereas [10] uses the one marginalized over the primordial CMB fields. This approach is also used by the MUSE framework [15], and was employed by the latest SPT-3G cosmological analysis [16]. The marginal and joint approaches are related in a precise manner: this second approach produces the same CMB and lensing maps as the first, were the first method to neglect all sources of mean fields, as discussed in Ref. [31], in particular, the delensed noise mean field. Hence our results are also relevant for the MUSE approach, since the maps produced by this method will contain this mean field.³

We then test the quantitative impact of this mean field on temperature-based reconstructions in Secs. III B to III D. Finally in Section III E, we investigate the impact of the delensed noise mean field on the delensing of the B modes. In this Section we work with the FFP10 cosmology.⁴ Note that we work at fixed cosmology, i.e., we do not evaluate the impact of using a nonfiducial cosmology in the simulations. In [12,13], we found that the impact of the cosmology used for the lensing reconstruction was negligible for the iterative lensing normalization, and the lensing biases ($N_L^{(0)}$ and $N_L^{(1)}$) can be mitigated by using a realization-dependent debiaser. In an end-to-end cosmological analysis like in [13], we can then include a first order correction on the response and on the biases with respect to the sampled lensing spectrum and CMB spectra, following [32]. Additionally, in [5], we found that the fiducial cosmology does not bias the recovered tensor to scalar ratio r from B -mode delensing. Thus we do not expect the delensed noise mean field to be impacted by a nonfiducial cosmology used in the simulations.

A. MAP estimator and delensed noise mean field

The key ingredient of likelihood-based reconstruction is a good model for the CMB covariance matrix conditioned on the deflections. Let the observed CMB Stokes fields be X^{dat} , and Cov_α the covariance. The simplest covariance matrix models will be of the form

$$\text{Cov}_\alpha \equiv \mathcal{B} \mathcal{D}_\alpha \mathcal{Y} C^{\text{unl}} \mathcal{Y}^\dagger \mathcal{D}_\alpha^\dagger \mathcal{B}^\dagger + N, \quad (3.1)$$

where $\mathcal{D}_\alpha$ is the lensing remapping operation ($X^{\text{len}} = \mathcal{D}_\alpha X^{\text{unl}}$, where X^{unl} is the unlensed CMB), $\mathcal{B}$ is the beam

³In the MUSE framework, a prior on the mean convergence value is introduced in order to mitigate the impact of the mean field [15].

⁴https://github.com/carronj/plancklens/blob/master/plancklens/data/cls/FFP10_wdipole_params.ini

of the experiment, $\mathcal{Y}$ is compact notation for the spherical harmonic transforms that map spherical harmonics coefficients of the unlensed CMB to the real space maps, C^{unl} is the covariance matrix of the unlensed CMB, which is diagonal in harmonic space, and N the noise covariance matrix.

The deflection field posterior is then

$$-2 \ln \mathcal{P}^{\text{M}}(\alpha | X^{\text{dat}}) = X^{\text{dat}\dagger} \text{Cov}_\alpha^{-1} X^{\text{dat}} + \ln \det \text{Cov}_\alpha + \kappa \cdot C^{\kappa\kappa, -1} \kappa. \quad (3.2)$$

The term

$$\kappa \cdot C^{\kappa\kappa, -1} \kappa \equiv \sum_{LM} \frac{|\kappa_{LM}|^2}{C_L^{\kappa\kappa}} \quad (3.3)$$

is a Gaussian prior on the lensing map, which plays no role in this discussion. The likelihood gradient at nonzero deflection field α contains the information on the residual lensing signal. Explicitly, one can define [11]

$$\frac{\delta}{\delta \alpha_a(\hat{n})} [\ln \mathcal{P}^{\text{M}}(\alpha | X^{\text{dat}})] \equiv {}_a g_\alpha^{\text{QD}}(\hat{n}) - {}_a g_\alpha^{\text{MF}}(\hat{n}) + {}_a g_\alpha^{\text{Pri}}(\mathbf{n}). \quad (3.4)$$

The terms g_α^{QD} and g_α^{MF} are the piece quadratic in the data and the mean field, respectively, and the subscript a denotes one of the two coordinates directions on the sphere. The quadratic gradient is the piece that tries to capture the residual lensing signal, beyond that represented by α , the current estimate of the deflection field. For example, $g_{\alpha=0}^{\text{QD}}$ is quite literally the standard quadratic estimator [21], prior to normalization, in the generalized minimum variance form discussed by [28]. For nonzero α , the form of g_α^{QD} is only slightly different, with the differences accounting for

- (1) first constructing the delensed CMB (“ $\hat{X}_\alpha^{\text{WF}}$ ”), under the assumption that α and the other ingredients of the likelihood model are the truth and
- (2) building then a quadratic estimator from these partially delensed maps, which will capture residual anisotropies.

More precisely, the quadratic piece may be written in the following form (see [31]),

$$g_\alpha^{\text{QD}}(\hat{n}) = |A_\alpha|(\hat{n}) [\mathcal{D}_\alpha g_\alpha^{\text{QE,del}}](\hat{n}), \quad (3.5)$$

which is a useful representation in order to understand the mean field: in this equation, $g_\alpha^{\text{QE,del}}$ can be unambiguously identified to a standard quadratic estimator on delensed maps, as defined in [31]. This quadratic estimator is of the “optimal filtering” type, (2.5), which includes delensed noise anisotropies in the filters. The prefactor $|A_\alpha| = |d^2 \hat{n}' / d^2 \hat{n}| \sim 1 - 2\kappa$ is the magnification matrix determinant

of the sphere remapping $\hat{n} \rightarrow \hat{n}'$ induced by α . The logic behind the factors $|A_\alpha|(\hat{n}) \mathcal{D}_\alpha$ in Eq. (3.5) is this: $g_\alpha^{\text{QE,del}}$ captures residual lensing on the delensed maps, hence it probes residual deflections at the *unlensed* positions, not the observed ones. These factors, which combine the remapping operation together with the distortions of the coordinates from unlensed to observed space, are precisely the ones obtained by the chain rule that connect the likelihood gradients with respect to the deflection field at the unlensed and lensed positions.

The second term in Eq. (3.4) corresponds to a mean field, it is the gradient of the log-determinant of the covariance. It is independent of the data, by definition. By noting that the first variation of the log-likelihood should be zero on average, we can write the mean-field gradient as

$$g_\alpha^{\text{MF}} = \langle g_\alpha^{\text{QD}} \rangle_{\text{fixed } \alpha} \quad (3.6)$$

where the average occurs over realizations of the data model (3.1) conditioned on α . According to Eq. (3.5), this is proportional to a mean field of a standard QE. It follows that

- (1) in the absence of anisotropies other than lensing, since the QE mean-field $\langle g_{\alpha=0}^{\text{QE,del}} \rangle$ vanishes, we can simply drop the prefactor when working to linear order in α , with the result

$$g_\alpha^{\text{MF}} \sim \langle g_\alpha^{\text{QE,del}} \rangle_{\text{fixed } \alpha}. \quad (3.7)$$

This shows that the mean-field linear (perturbative) response $\mathcal{R}_L^{\text{MF}}$ to α will be that of a QE with anisotropic filtering, defined just as we did in the previous section in Eq. (2.6).

- (2) In practice, the zeroth-order QE mean field will never be zero. This part then just gets remapped, and modulated by the magnification.

We note that the perturbative prediction of the mean field from Eq. (2.6) is equivalent to modifying the prior on the lensing field in Eq. (3.3). Indeed, the perturbative mean field can be absorbed in the prior, which becomes

$$\kappa \cdot C^{\kappa\kappa, -1} \kappa \rightarrow \kappa \cdot (C^{\kappa\kappa, -1} - \mathcal{R}^{\kappa\kappa, \text{MF}}) \kappa. \quad (3.8)$$

As we will discuss in Sec. III D, the predicted normalization of the MAP estimator is a Wiener filter driven by the prior. Thus neglecting the delensing mean field will change the normalization of the MAP estimator.

We show in Fig. 4 an estimate of the delensed mean field at the first iteration of the MAP estimator. Thus the deflection field α on this figure is obtained from the quadratic estimator. We show on the left panel the mean-field g_α^{MF} from the perturbative expression of Eq. (2.6), and on the central panel the mean field estimated from averaging over 20 simulations at fixed deflection field α . We see that the perturbative mean field is in good visual agreement with the one estimated from




FIG. 4. Estimation of the delensed noise mean-field κ^{MF} in real space, at the first iteration of the MAP estimator from temperature. Simulations are with CMB-S4 like noise level. The left panel shows the predicted delensed noise mean field at first order in the deflection field. The central panel shows the delensed noise mean field estimated from a set of 20 simulations at fixed deflection field. The right panel shows the negative of the Wiener filtered true convergence field of the simulation. Patches are 200×200 pixels, with 1.5 arcmin per pixel, cutouts from full-sky simulations.

the simulations. The right panel shows that the delensed noise mean field is indeed correlated with the true lensing field, as expected.

B. Mean-field tests

When reconstructing the lensing field with the MAP estimator, the mean-field gradient term should in principle be estimated at each iteration of the algorithm. However, performing Monte-Carlo estimation of this mean field at each iteration is computationally expensive and induces residual Monte-Carlo noise in the gradient. Hence there is value in exploring other possibilities to treat this term. We consider different schemes to see the impact of the MF gradient term in the iterations.

- (i) Neglecting the MF. We neglect the MF and perform the lensing iterations assuming $g_{\alpha}^{\text{MF}} = 0$. In practice, this corresponds to truncating the likelihood gradient in Eq. (3.4) during the MAP search. As such the lensing potential at convergence might not correspond to the true maximum of the posterior. However it has the advantage of being computationally cheap. This also matches the implementation of the CMB-joint MAP estimator used in MUSE, which does not include the MF term.
- (ii) Simulated MF. We estimate the MF using the relation $g_{\alpha}^{\text{MF}} = \langle g_{\alpha}^{\text{QD}} \rangle$. In practice, for each iteration in the MAP search, we estimate the average of the quadratic gradient term g_{α}^{QD} over a set of 20 simulations, lensed by the estimated deflection field α^{MAP} . We decrease the variance of this MF estimate by using the tricks of the Appendix B2 of [11]. Specifically, we modify the quadratic gradient term to operate on maps with unit variance instead of the CMB power spectrum variance, while preserving the expected value. Additionally, we further reduce variance by subtracting an MF estimate derived from the same set of unlensed CMB simulations. To remove any potential bias, at each iteration we

use a different set of simulations to estimate the mean field.

- (iii) Perturbative MF. We estimate the delensed noise mean field using the perturbative approach described in Sec. II. The MF estimate is thus given by $g_{\alpha, LM}^{\text{MF}} = \mathcal{R}_L^{\kappa, \text{MF}} \hat{\kappa}_{LM}^{\text{MAP}}$, where $\mathcal{R}_L^{\kappa, \text{MF}}$ is obtained from Eq. (2.8) using the noise spectra in the CMB response and delensed CMB spectra in the QE weights.

We consider three different CMB configurations, reproducing the Planck, Simons Observatory, and CMB-S4 white noise. The corresponding beam, noise levels, and CMB scale cuts are described in Table I. We generate ideal full-sky simulations, without any sources of anisotropies other than the lensing field itself. In practice we generate a Gaussian unlensed CMB temperature map, lens it with a Gaussian lensing field, and add a realization of the noise. We reconstruct the lensing field on the full sky with the MAP estimator as described in detail in [5]. We rely only on the temperature map to perform the lensing reconstructions.

C. Impact on lensing reconstruction

For each experimental setting, we estimate the correlation coefficient between the reconstructed MAP lensing field $\hat{\phi}^{\text{it}}$ and the simulation input lensing field ϕ^{in}

$$\rho_L^{\text{MAP}} = \frac{C_L^{\hat{\phi}^{\text{it}}, \phi^{\text{in}}}}{\sqrt{C_L^{\hat{\phi}^{\text{it}}, \hat{\phi}^{\text{it}}} C_L^{\phi^{\text{in}}, \phi^{\text{in}}}}}. \quad (3.9)$$

TABLE I. Experimental configuration of our simulations.

Experiment	Planck	Simons Obs.	CMB-S4
Noise ($\mu\text{K} \cdot \text{arcmin}$)	35	3	1
Beam FWHM (arcmin)	5	3	1
ℓ_{min}	100	40	30
ℓ_{max}	2048	4000	4000

The predicted correlation coefficient is given by

$$(\rho_L^{\text{fid}})^2 = \frac{C_L^{\phi\phi,\text{fid}}}{C_L^{\phi\phi,\text{fid}} + N_L^{(0)} + N_L^{(1)}}, \quad (3.10)$$

where $N_L^{(0)}$ and $N_L^{(1)}$ are the iterative lensing reconstruction biases, estimated from the fiducial procedure of [12,13].

Figure 5 compares the lensing correlation coefficient for our different MF scenarios, after ten iterations of the MAP estimator. The predicted correlation coefficient is shown in the upper panel, for the MAP (teal) and QE (golden) reconstructions. The lower panels show the correlation coefficient after ten iterations of our simulations, normalized by the fiducial correlation coefficients, for Planck, SO, and CMB-S4 noise levels from top to bottom. We bin the correlation coefficient in L -bins of size $\Delta L = 150$, showed with offset for clarity. Error bars are the standard error on the mean of the correlation coefficient in each bin.

The agreement with the predicted correlation coefficient is at the 2% level. The prediction of the $N_L^{(0)}$ and $N_L^{(1)}$ iterative biases are not perfect, as discussed in [12,13], since they are obtained from a heuristic approach, and this is probably the main source of discrepancy. The bias in the QE correlation coefficient for CMB-S4 is also at the 2% level, this might be due to the fact that we work on the full sky, and the $N_L^{(1)}$ prediction is obtained in the flat-sky approximation.

We see that using a different recipe for the delensed noise mean field, either perturbative or from the simulations, or completely neglecting it, has no impact on the correlation coefficient at all scales.

We note that for the lowest multipole bin, the perturbative MF has a larger error bar than the simulated MF. This can be due to the fact that the perturbative approximation breaks down, as the approximation is only valid at first order in the deflection field. However, this does not impact the correlation coefficient itself.

On top of these tests, we also looked at neglecting the delensed noise mean field in all iterations but the last one, or using a simulated MF with a filtering of the small scales (which contain Monte-Carlo noise). We found that these methods do not improve over the simple case of neglecting the mean field in all iterations.

D. Power spectrum normalization

We now consider the impact of the delensed noise mean field on the normalization of the MAP lensing potential.

We recall that our best prediction for the MAP lensing power spectrum normalization is given by a Wiener filter

$$\mathcal{W}_L = \frac{C_L^{\phi\phi,\text{fid}}}{C_L^{\phi\phi,\text{fid}} + N_L^{(0)}}, \quad (3.11)$$




FIG. 5. Upper panel: Predicted correlation coefficients for the temperature only MAP (teal) and QE (golden) reconstruction for Planck, SO, and CMB-S4-like noise levels, respectively, in dotted, dash-dotted, and plain lines. Lower panels: Correlation coefficient after ten iterations of our lensing reconstruction divided by the fiducial correlation coefficient, for Planck, SO, and CMB-S4 noise levels from top to bottom. Note that the range of L value is not the same in each panel. In teal we show the case without mean-field subtraction, in pink it is using the perturbative mean-field prediction and in purple it is with the MF estimated from simulations. The QE case is shown in gold for comparison. We see that for all recipes we find very similar correlation coefficients: neglecting the delensed noise mean field has no impact on the quality of the lensing potential reconstruction. We also see that the predicted correlation coefficients are accurate at $\sim 2\%$.

where $C_L^{\phi\phi,\text{fid}}$ is the fiducial lensing power spectrum and $N_L^{(0)}$ is the iterative $N_L^{(0)}$ bias. However we showed in [12,13] that this prediction is not accurate, and has to be corrected from a set of simulations. In practice for the polarization




FIG. 6. Isotropic normalization of the MAP lensing potential, for the Planck, SO, and S4 noise levels from top to bottom. In teal we show the case without mean-field subtraction, in purple it is with the mean field estimated from simulations, and in pink it is when using the perturbative mean-field prediction. The black plain line is the fiducial MAP response and the dashed black line is the fiducial MAP response when we neglect the delensed noise mean field in the likelihood.

estimator we found that the prediction is off by around 3%. Crucially, this correction is independent of the CMB and noise power spectra contained in the map: it only depends on the fiducial ingredients used for the MAP reconstruction.

We show in Fig. 6 the effective normalization for the MAP lensing field with the temperature only estimator. This effective normalization is estimated from the cross-correlation between the MAP lensing field $\hat{\phi}^{it}$ after ten iterations, with the true lensing field of the simulation ϕ^{in}

$$\hat{\mathcal{W}}_L = \frac{C_L^{\hat{\phi}^{it}, \phi^{in}}}{C_L^{\phi^{in}, \phi^{in}}}. \quad (3.12)$$

We see that when neglecting the delensed noise mean field in the iterations, the effective normalization is decreased.

For the Simons observatory, the effective normalization is decreased by about 15% for $L < 500$, and for CMB-S4 by about 20%. This is more important than the $\sim 3\%$ deviation we found on the polarization only estimator in [13].

The expression of the normalization as a Wiener filter is an approximation of the true response of the MAP estimator. This response is given by

$$\mathcal{W} = \left[\frac{1}{C_L^{\phi\phi, \text{fid}}} + \mathcal{H} \right]^{-1} \mathcal{R}, \quad (3.13)$$

where $\mathcal{R} = \frac{\delta g^{\text{QD}}}{\delta X^{\text{dat}}} \frac{\delta X^{\text{dat}}}{\delta \phi}$ is the response of the quadratic gradient to the true lensing potential, analog of the standard QE response, and $\mathcal{H} = \frac{\delta^2 \ln \mathcal{L}}{\delta \phi^i \delta \phi^j}$ is the curvature of the likelihood estimated in ϕ^{MAP} . In the signal dominated regime we can approximate $\mathcal{H} \simeq \mathcal{R}$, which gives back the expression of the Wiener filter in Eq. (3.11).

When ignoring the delensed noise mean field in the likelihood, we can approximate the likelihood curvature as $\mathcal{H} \simeq \mathcal{R} - \mathcal{R}^{\text{MF}}$, where we have subtracted the delensed noise mean-field response $\mathcal{R}^{\text{MF}}$. This gives the following fiducial expression for the normalization

$$\mathcal{W}_L = \frac{C_L^{\phi\phi, \text{fid}}}{C_L^{\phi\phi, \text{fid}} + N_L^{(0)} - \frac{\mathcal{R}^{\text{MF}}}{\mathcal{R}_L} C_L^{\phi\phi, \text{fid}}}. \quad (3.14)$$

This prediction is shown as the dashed black lines in Fig. 6. We see that in all experimental configurations, this prediction is a good approximation of the effective normalization obtained when neglecting the delensed noise mean field in the iterations. The prediction for the response without mean field is accurate at $\sim 8\%$ for CMB-S4 at $L = 500$.

E. Delensing

Delensing the polarization B modes allows one to put tight constraints on the tensor to scalar ratio. The delensing can be performed using a template of the lensed B modes, obtained by combining an estimate of the lensing potential with the observed E modes of polarization [18]. Here we assume that the lensing potential is reconstructed from the MAP temperature only estimator. In reality, one would use as well the polarization MAP estimator, or both temperature and polarization together. But we now wish to test the impact of the delensed noise mean field on the obtained B mode template, which is better seen in the temperature only estimator.

We construct the B mode template perturbatively, using the MAP lensing potential (obtained after ten iterations), and the observed (lensed) E modes, which have been Wiener filtered to remove the noisy small scale modes. We do not renormalize the lensing potential ϕ^{it} to construct the




FIG. 7. Residual B -modes power spectra after subtracting B -modes templates estimated from temperature only estimators, for a CMB-S4 configuration. The gold triangles show the delensing when the lensing template is obtained with a quadratic estimator. We show different scenarios for the MAP search: neglecting the mean field (teal), mean field estimated from Monte-Carlo simulations (purple), or with a perturbative prediction (pink). We binned the spectra in ten multipole bins between $\ell_{\min} = 30$ and $\ell_{\max} = 200$, the bin centers are offset for clarity. The black line shows the lensed B modes power spectrum for reference.

lensing template. Indeed, we expect the MAP lensing potential to be optimally Wiener filtered.

We show in Fig. 7 the residual B modes power spectrum when delensing with the lensing template $B^{\text{del}} = B^{\text{dat}} - \hat{B}^{\text{temp}}$. We see that neglecting the delensed noise mean field in the MAP iterations does not impact the delensing of the B modes. The efficiency of the delensing is not improved when considering the mean field estimated from simulations or perturbatively at each iteration. The negligible impact of the delensed noise mean field on delensing might be attributed to the fact that the lensed B modes are predominantly generated by the small-scale lensing potential, as illustrated in Figs. 11 of [32] and 1 of [5]. In contrast, the delensed noise mean field primarily affects the large scales of the lensing potential. Furthermore, as shown in Fig. 5, neglecting the delensed noise mean field still results in an optimal reconstruction of the lensing potential field.

IV. CONCLUSION

Delensing the CMB removes anisotropies from the CMB signal but also creates new anisotropies in the noise maps. For any quadratic estimator applied on delensed CMB maps, these anisotropies will appear as a *delensed noise mean field*. This effect can, in principle, bias the field of interest reconstructed from these delensed CMB maps.

We have shown that the main effect is a modulation of the noise variance maps, more relevant in temperature than in polarization, and that to leading order this delensed noise

mean field traces very directly the large-scale lensing convergence map κ . We derived an analytic expression for the response of this delensed noise mean field at first order in κ , with the help of calculations standard in the context of quadratic estimators.

We then studied the impact of this delensed noise mean field within the framework of the MAP estimator and found that its noticeable effect is to change the normalization of the reconstructed lensing field.

Since at perturbative order the delensed noise mean field is proportional to the lensing convergence map κ , it can be absorbed into the Gaussian prior of the MAP estimator. This also explains why neglecting it mostly impacts the normalization of the MAP estimator, since the Gaussian prior is driving the Wiener filter normalization. If the mean-field response is indeed isotropic and perturbative, then the corresponding curvature matrix of the likelihood is effectively constant and diagonal. This means it would not induce mode mixing, but merely rescales the normalization of the reconstructed field.

In Eq. (3.14) we have obtained a good analytical estimate of the MAP normalization in the case that the mean field is neglected during the lensing reconstruction. A more complete understanding of the normalization would require additional analytical insight into the dense curvature matrix of the likelihood. This remains a formidable task.

Provided that the normalization is consistently estimated with a set of simulations, it is safe to neglect the delensed noise mean field in the MAP search. Indeed, in particular for cross-correlation analysis [33–35], one should be careful that whenever the noise properties of the survey vary across the sky, the normalization has to be recomputed for the footprint of the survey used in the analysis. However we found that the delensed noise mean field does not worsen the correlation coefficient of the reconstructed lensing field, and it does not impact the B mode’s delensing. Neglecting the delensed noise mean field allows for faster MAP search, as it saves the computational time needed to estimate the mean field in between iterations.

Therefore, for most cosmological analysis there appears to be little impact of the delensed noise mean field, provided the change in normalization of the estimator is correctly accounted for. However, failure to do so in cross-correlations of CMB lensing with large-scale structure could lead to significant biases of about 15% to 20% for a Simons-Observatory-like or CMB-S4-like experiment, respectively.

CMB foregrounds such as SZ, CIB, and radio sources are also a source of anisotropy in the CMB maps, and they are highly non-Gaussian. They can bias the CMB lensing reconstruction, and since they are correlated to the density field, they could in principle be correlated with the delensed noise mean field. Mitigation techniques such as the bias hardening [26,36] can be used to reduce these foreground biases. However, a more complete study of the interplay of

foregrounds and the delensed noise mean field is left for future work.

Finally, our analysis showed that the delensed noise mean field should not be neglected when applying a quadratic estimator sensitive to modulationlike signals on delensed CMB maps, such as in searches for patchy reionization [23,26].

Some computations used resources provided through the STFC DiRAC Cosmos Consortium and hosted at the Cambridge Service for Data Driven Discovery (CSD3).

ACKNOWLEDGMENTS

L.L. acknowledges support from a SNSF Postdoc Mobility fellowship (Grant No. P500PT_21795), and from the São Paulo Research Foundation (FAPESP), Brasil Process No. 2023/08560-7. J. C. acknowledges support from a SNSF Eccellenza Professorial Fellowship (No. 186879). This work was supported by a grant from the Swiss National Supercomputing Centre (CSCS) under Project No. sm80.

DATA AVAILABILITY

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

APPENDIX: QUADRATIC ESTIMATORS NOISE MEAN FIELDS

In this appendix we discuss further the leading order full-sky MAP mean field, its connection to more standard noise mean fields, and how it can be computed fairly simply with already available lensing codes.

According to the arguments in the main text Sec. III A, we can obtain the leading contribution from a quadratic estimator. This quadratic estimator is obtained in the standard manner, with the understanding that the filtering is anisotropic (α is present in the filtering matrix, since this is known source of anisotropy in the lensing reconstruction process beyond the quadratic estimator).

By definition, mean-field calculations must be performed under the assumption that the input maps match precisely the likelihood model. Using this, we get the very compact formula

$$\langle_{\pm 1} \hat{g}_{\alpha}^{\text{QE,del}}(\hat{n}) \rangle = - \sum_{l_1 m_1 l_2 m_2} \mathcal{W}_{\alpha}^{l_1 m_1 l_2 m_2} \delta_{\pm 0}^{\pm} Y_{l_1 m_1}(\hat{n}) Y_{l_2 m_2}^*(\hat{n}) \quad (\text{A1})$$

where $\hat{g}_{\alpha}^{\text{QE,del}}$ is introduced in Eq. (3.5) as the standard quadratic estimator on delensed maps, with anisotropic filtering depending on the deflection field α ; $\mathcal{W}_{\alpha}$ is the Wiener-filtering matrix; $\mathcal{W}_{\alpha} = (C^{-1} + N_{\alpha}^{-1})^{-1} N_{\alpha}^{-1}$, which may also be written formally as $C(C + N_{\alpha})^{-1}$ and $Y_{lm}(\hat{n})$ are the spherical harmonics. This equation may look a bit mysterious but can be given some interpretation: consider

first introducing two different locations $\hat{n}_1$ and $\hat{n}_2$ in the respective harmonics (instead of the identical $\hat{n}$). After doing so, the object on the right-hand side in this equation can be seen as a matrix with indices $\hat{n}_1, \hat{n}_2$. This matrix represents the real-space operation of producing the most probable (in the rigorous, Bayesian sense—this is what a Wiener filter is doing) estimate of the gradient of a map at $\hat{n}_1$, from the input map at $\hat{n}_2$, and the mean field is the diagonal ($\hat{n}_1 = \hat{n}_2 = \hat{n}$) of this map-to-gradient matrix. This comes about because the defining feature of (perturbative) lensing is to introduce in the temperature map at any location $\hat{n}$ a term proportional to its gradient at this very same location. In the absence of other anisotropies (in our full-sky case, evaluating the likelihood model at vanishing α), no other mechanism does that. The value of the map and its gradient at identical locations are statistically independent: knowledge of the map at $\hat{n}$ does not help estimating the gradient there, so that this average vanishes. If there are known anisotropies in the map (known in the sense of being implemented in the likelihood model), optimal recovery of the gradient may in general involve to some extent the map at the same location. Hence, the map and its recovered gradient are now related, but this is not lensing and must be subtracted.

It is useful to compare this formula to that obtained in the case of more standard quadratic estimators. To do this, to avoid confusion with lensing-induced anisotropies, let β denote now any source of anisotropy. Quadratic estimators are very flexible and used in different variations, but the relevant aspects can be summarized as a two step procedure: 1) inverse-covariance weighting of CMB maps using some covariance model 2) application of weights specific to the targeted source of anisotropy to these maps. Hence it is also useful to distinguish explicitly the true covariance matrix of the CMB in the presence of the anisotropy [let us call it $(C^{\text{dat}})_{\beta}$] from the covariance matrix model used in step 1). Let us call the filtering matrix (F_{β}) . For optimal filtering, this matrix is the inverse of the data covariance, but in many cases the anisotropy is neglected in the filters, in which case the filtering $F_{\beta} \equiv F_0$ is isotropic in harmonic space. For lensing, step 2) is what produces the Wiener-filtered gradient from the inverse-covariance filtered maps. This involves using a set of isotropic fiducial spectra, which we call C_0^w , to possibly distinguish it from the CMB spectra. With this notation in place, the general lensing-QE response to β is given by the same formula Eq. (A1), but replacing the matrix $\mathcal{W}_{\alpha}$ with

$$\mathcal{W}_{\alpha} \rightarrow C_0^w F_{\beta} C_{\beta}^{\text{dat}} F_{\beta}. \quad (\text{A2})$$

We can distinguish at least three cases of interest.

- (1) The anisotropy lies in the CMB signal, and the QE uses isotropic filters. In this case, (A2) becomes

$$C_0^w F_0 C_{(1),\beta} F_0 \quad (\text{A3})$$

where $C_{(1)\beta}$ is the leading order change to the CMB spectra (a dense matrix).

- (2) The anisotropy is now present not in the signal but in the noise CMB maps, but just as for point 1) it is neglected in the filters. In this case the relevant matrix is

$$C_0^w F_0 N_{(1)\beta} F_0, \quad (\text{A4})$$

i.e., this just amounts to trivially replacing one set of CMB spectra with the noise spectra.

- (3) The anisotropy is present in the noise CMB map, and is implemented in the filters F . For weak anisotropies, the effect of anisotropic filtering is that the response changes sign: explicitly, to leading perturbative order,

$F_\beta \sim F_0 - F_0 N_{(1)\beta} F_0$, and plugging this in (A2) one gets

$$-C_0^w F_0 N_{(1)\beta} F_0. \quad (\text{A5})$$

Hence, comparing points 1) and 3) demonstrates very explicitly a recipe to compute the leading MAP mean-field response to α as follows: just use a standard lensing response calculation code, but replace one set of CMB spectra in the numerator (the one entering the sky response to lensing $C_{(1)\alpha}$) by the noise spectra. The other spectra C_0^w and the filtering spectra remain unchanged. There is no explicit change in sign, since for the MAP mean field in Eq. (A5) the anisotropy β is tracing $-\alpha$, canceling the sign in front, whereas the standard lensing response (A2) has $\beta = +\alpha$.

-
- [1] Antony Lewis and Anthony Challinor, Weak gravitational lensing of the CMB, *Phys. Rep.* **429**, 1 (2006).
 - [2] M. Tristram *et al.*, Improved limits on the tensor-to-scalar ratio using BICEP and Planck data, *Phys. Rev. D* **105**, 083524 (2022).
 - [3] P.A.R. Ade *et al.* (BICEP and Keck Collaborations), Improved constraints on primordial gravitational waves using Planck, WMAP, and BICEP/Keck observations through the 2018 observing season, *Phys. Rev. Lett.* **127**, 151301 (2021).
 - [4] Julien Carron, Antony Lewis, and Anthony Challinor, Internal delensing of Planck CMB temperature and polarization, *J. Cosmol. Astropart. Phys.* **05** (2017) 035.
 - [5] 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).
 - [6] Toshiya Namikawa *et al.*, Simons observatory: Constraining inflationary gravitational waves with multitracer B -mode delensing, *Phys. Rev. D* **105**, 023511 (2022).
 - [7] Emilie Hertig *et al.* (Simons Observatory Collaboration), The Simons Observatory: Combining cross-spectral foreground cleaning with multitracer B -mode delensing for improved constraints on inflation, *Phys. Rev. D* **110**, 043532 (2024).
 - [8] Daniel Green, Joel Meyers, and Alexander van Engelen, CMB delensing beyond the B modes, *J. Cosmol. Astropart. Phys.* **12** (2017) 005.
 - [9] 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.
 - [10] Christopher M. Hirata and Uros Seljak, Reconstruction of lensing from the cosmic microwave background polarization, *Phys. Rev. D* **68**, 083002 (2003).
 - [11] Julien Carron and Antony Lewis, Maximum *a posteriori* CMB lensing reconstruction, *Phys. Rev. D* **96**, 063510 (2017).
 - [12] Louis Legrand and Julien Carron, Robust and efficient CMB lensing power spectrum from polarization surveys, *Phys. Rev. D* **108**, 103516 (2023).
 - [13] Louis Legrand and Julien Carron, Lensing power spectrum of the cosmic microwave background with deep polarization experiments, *Phys. Rev. D* **105**, 123519 (2022).
 - [14] Marius Millea, Ethan Anderes, and Benjamin D. Wandelt, Bayesian delensing of CMB temperature and polarization, *Phys. Rev. D* **100**, 023509 (2019).
 - [15] Marius Millea and Uros Seljak, Marginal unbiased score expansion and application to CMB lensing, *Phys. Rev. D* **105**, 103531 (2022).
 - [16] F. Ge *et al.* (SPT-3G Collaboration), Cosmology from CMB lensing and delensed EE power spectra using 2019–2020 SPT-3G polarization data, *Phys. Rev. D* **111**, 083534 (2025).
 - [17] Patricia Larsen, Anthony Challinor, Blake D. Sherwin, and Daisy Mak, Demonstration of cosmic microwave background delensing using the cosmic infrared background, *Phys. Rev. Lett.* **117**, 151102 (2016).
 - [18] Antón Baleato Lizancos, Anthony Challinor, and Julien Carron, Limitations of CMB B -mode template delensing, *Phys. Rev. D* **103**, 023518 (2021).
 - [19] Y. Akrami *et al.* (Planck Collaboration), Planck 2018 results. IV. Diffuse component separation, *Astron. Astrophys.* **641**, A4 (2020).
 - [20] Julien Carron, Mark Mirmelstein, and Antony Lewis, CMB lensing from Planck PR4 maps, *J. Cosmol. Astropart. Phys.* **09** (2022) 039.
 - [21] Takemi Okamoto and Wayne Hu, CMB lensing reconstruction on the full sky, *Phys. Rev. D* **67**, 083002 (2003).

- [22] P. A. R. Ade *et al.* (Planck Collaboration), Planck 2015 results. XV. Gravitational lensing, *Astron. Astrophys.* **594**, A15 (2016).
- [23] Cora Dvorkin and Kendrick M. Smith, Reconstructing patchy reionization from the cosmic microwave background, *Phys. Rev. D* **79**, 043003 (2009).
- [24] Toshiya Namikawa, CMB mode coupling with isotropic polarization rotation, *Mon. Not. R. Astron. Soc.* **506**, 1250 (2021).
- [25] P. A. R. Ade *et al.* (BICEP/Keck Collaboration), BICEP/Keck. XVII. Line-of-sight distortion analysis: Estimates of gravitational lensing, anisotropic cosmic birefringence, patchy reionization, and systematic errors, *Astrophys. J.* **949**, 43 (2023).
- [26] Omar Darwish, Optimal joint reconstruction from CMB observations: Application to cosmic birefringence, patchy reionization and CMB lensing, *Phys. Rev. D* **112**, 023514 (2025).
- [27] Antony Lewis, Anthony Challinor, and Duncan Hanson, The shape of the CMB lensing bispectrum, *J. Cosmol. Astropart. Phys.* 03 (2011) 018.
- [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] Martin Bucher, Carla Sofia Carvalho, Kavilan Moodley, and Mathieu Remazeilles, CMB lensing reconstruction in real space, *Phys. Rev. D* **85**, 043016 (2012).
- [30] Julien Carron and Antony Lewis, Spherical bispectrum expansion and quadratic estimators, *J. Cosmol. Astropart. Phys.* 07 (2024) 067.
- [31] Julien Carron, On likelihood-based analysis of the gravitationally (de)lensed CMB, [arXiv:2502.02399](https://arxiv.org/abs/2502.02399).
- [32] N. Aghanim *et al.*, Planck 2018 results. VIII. Gravitational lensing, *Astron. Astrophys.* **641**, A8 (2020).
- [33] Aurelien Benoit-Levy, Typhaine Dechelette, Karim Benabed, Jean-Francois Cardoso, Duncan Hanson, and Simon Prunet, Full-sky CMB lensing reconstruction in presence of sky-cuts, *Astron. Astrophys.* **555**, A37 (2013).
- [34] Julien Carron, Real-world CMB lensing quadratic estimator power spectrum response, *J. Cosmol. Astropart. Phys.* 02 (2023) 057.
- [35] Gerrit S. Farren *et al.* (ACT Collaboration), The Atacama Cosmology Telescope: Cosmology from cross-correlations of unWISE Galaxies and ACT DR6 CMB lensing, *Astrophys. J.* **966**, 157 (2024).
- [36] Toshiya Namikawa, Duncan Hanson, and Ryuichi Takahashi, Bias-hardened CMB lensing, *Mon. Not. R. Astron. Soc.* **431**, 609 (2013).