



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Optimal multifrequency weighting for CMB lensing

Noah Sailer, Emmanuel Schaan, Simone Ferraro, Omar Darwish, and Blake Sherwin

Phys. Rev. D **104**, 123514 — Published 6 December 2021

DOI: [10.1103/PhysRevD.104.123514](https://doi.org/10.1103/PhysRevD.104.123514)

Optimal multi-frequency weighting for CMB lensing

Noah Sailer,^{1,2,*} Emmanuel Schaan,^{2,1} Simone Ferraro,^{2,1} Omar Darwish,³ and Blake Sherwin^{3,4}

¹*Berkeley Center for Cosmological Physics, Department of Physics,
University of California, Berkeley, CA 94720, USA*

²*Lawrence Berkeley National Laboratory, One Cyclotron Road, Berkeley, CA 94720, USA*

³*Department of Applied Mathematics and Theoretical Physics,
University of Cambridge, Wilberforce Road, Cambridge CB3 0WA, UK*

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

Extragalactic foregrounds in Cosmic Microwave Background (CMB) temperature maps lead to significant biases in CMB lensing reconstruction if not properly accounted for. Combinations of multi-frequency data have been used to minimize the overall map variance (internal linear combination, or ILC), or specifically null a given foreground, but these are not tailored to CMB lensing. In this paper, we derive an optimal multi-frequency combination to jointly minimize CMB lensing noise and bias. We focus on the standard lensing quadratic estimator, as well as the “shear-only” and source-hardened estimators, whose responses to foregrounds differ. We show that an optimal multi-frequency combination is a compromise between the ILC and joint deprojection, which nulls the thermal Sunyaev-Zel’dovich (tSZ) and Cosmic Infrared Background (CIB) contributions. In particular, for a Simons Observatory-like experiment with $\ell_{\max,T} = 3000$, we find that profile hardening alone (with the standard ILC) reduces the bias to the lensing power amplitude by 40%, at a 20% cost in noise, while the bias to the cross-correlation with a LSST-like sample is reduced by nearly an order of magnitude at a 10% noise cost, relative to the standard quadratic estimator. With a small amount of joint deprojection the bias to the profile hardened estimator can be further reduced to less than half the statistical uncertainty on the respective amplitudes, at a 20% and 5% noise cost for the auto- and cross-correlation respectively, relative to the profile hardened estimator with the standard ILC weights. Finally, we explore possible improvements with more aggressive masking and varying $\ell_{\max,T}$.

I. INTRODUCTION

CMB lensing [1–3] allows us to reconstruct the projected mass distribution all the way to the surface of last scattering, enabling cosmological inference (e.g. neutrino masses, dark matter and dark energy) using a very well-understood source redshift and redshift weighting. Extragalactic foregrounds, which are correlated with the lensing potential, are likely to be one of the main limiting factors [4–10] in temperature-based CMB lensing reconstruction, producing significant biases to the standard minimum variance quadratic estimator (QE, [11]), which in turn bias both the reconstructed lensing power spectrum and the cross-correlations with tracers of the large-scale structure. In this paper we consider two techniques to mitigate these biases: geometric-based methods that distinguish lensing from foregrounds through their different spatial symmetries (shear-only reconstruction [8] and bias-hardening [5, 10]), and multi-frequency based methods that exploit the different frequency scaling of the CMB and extragalactic foregrounds [6, 12]. In particular, we explore the *lensing-optimized* linear combinations of frequency maps to minimize a combination of bias and noise, for the standard QE, as well as for shear and bias-hardened estimators. Another promising method to control foreground biases to CMB lensing is the gradient-cleaned [6] or symmetrized estimator [13, 14]. We explore these methods in a companion paper [15]. The remainder of the paper is organized as follows: in §II we introduce

the standard, shear-only, and bias-hardened quadratic estimators; in §III we review multi-frequency methods for noise and bias mitigation; in §IV (V) we outline the calculation of the noise (bias) of a general quadratic estimator; in §VI we combine geometric- and frequency-based techniques to minimize a combination of lensing noise and bias; we conclude in §VII.

II. CMB LENSING QUADRATIC ESTIMATORS

In the presence of a foreground s_ℓ , the observed temperature covariance receives off-diagonal contributions from lensing and the foreground¹ (see Eqs. 3-7 in [10]):

$$\langle T_\ell T_{L-\ell} \rangle = f_{\ell,L-\ell}^\kappa \kappa_L + f_{\ell,L-\ell}^s s_L \quad (1)$$

to lowest order in κ and s . The standard quadratic estimator (QE, [11]) neglects the foreground term and seeks the minimum variance solution for κ_L in Eq. (1). By neglecting the foreground contribution, it can lead to a biased answer [8, 10].

On small scales where the foreground contributions are large, they are well approximated by a collection of unclustered halos or galaxies with azimuthally-symmetric profiles. In this regime, the shear-only κ estimator, which is insensitive to such isotropic terms has been shown to significantly reduce all foreground-induced biases [8].

¹ $T_\ell = T_\ell^{\text{CMB}} + s_\ell$, where T_ℓ^{CMB} is the *lensed* primary CMB, and the average is taken over realizations of the unlensed CMB at fixed κ_L and s_L .

* nsailer@berkeley.edu

Bias-hardening [5, 10, 16] seeks to measure the foreground contribution directly and “subtract off” the bias to the lensing convergence $\kappa_{\mathbf{L}}$. Since both $\kappa_{\mathbf{L}}$ and $s_{\mathbf{L}}$ produce off-diagonal covariances, the standard quadratic estimators for κ and s are biased [5, 10]:

$$\begin{pmatrix} \langle \hat{\kappa}_{\mathbf{L}} \rangle \\ \langle \hat{s}_{\mathbf{L}} \rangle \end{pmatrix} = \begin{pmatrix} 1 & N_{\mathbf{L}}^{\kappa} \mathcal{R}_{\mathbf{L}} \\ N_{\mathbf{L}}^s \mathcal{R}_{\mathbf{L}} & 1 \end{pmatrix} \begin{pmatrix} \kappa_{\mathbf{L}} \\ s_{\mathbf{L}} \end{pmatrix} \quad (2)$$

where the response $\mathcal{R}_{\mathbf{L}}$ is defined in [10]. A simple matrix inversion leads to unbiased estimates of the lensing and foreground:

$$\begin{pmatrix} \hat{\kappa}_{\mathbf{L}}^{\text{BH}} \\ \hat{s}_{\mathbf{L}}^{\text{BH}} \end{pmatrix} = \begin{pmatrix} 1 & N_{\mathbf{L}}^{\kappa} \mathcal{R}_{\mathbf{L}} \\ N_{\mathbf{L}}^s \mathcal{R}_{\mathbf{L}} & 1 \end{pmatrix}^{-1} \begin{pmatrix} \hat{\kappa}_{\mathbf{L}} \\ \hat{s}_{\mathbf{L}} \end{pmatrix}. \quad (3)$$

We assume that the foreground s can be approximated by a collection of sources with identical profiles u_{ℓ} [10]. Explicitly, we assume $s_{\ell} = \sum_i s_i u_{\ell}$, where s_i is the amplitude of the i 'th source. In this work we consider either hardening against point sources [5, 10, 16] (point source hardening, or PSH), where the profile is taken to be a delta function ($u_{\ell} = 1$), or against a tSZ-like profile [10] (profile hardening, or PH), where the profile is taken to be the square root of the tSZ power spectrum at 150 GHz.

III. INTERNAL LINEAR COMBINATION (ILC) METHODS

Let $(\mathbf{T}_{\ell})_i$ denote the observed map in the i -th frequency channel. We assume that \mathbf{T}_{ℓ} can be written in the form:

$$\mathbf{T}_{\ell} = \mathbf{1} T_{\ell}^{\text{CMB}} + \mathbf{n}_{\ell} + \sum_s \mathbf{A}_s s_{\ell}, \quad (4)$$

where $\mathbf{1} = (1, 1, \dots, 1)^T$ and T_{ℓ}^{CMB} is the *lensed* CMB. The “noise” \mathbf{n}_{ℓ} receives contributions from the detectors’ thermal fluctuations, the atmosphere, and galactic dust. The extragalactic foregrounds s_{ℓ} include tSZ, kSZ, CIB, and radio point sources. In Eq. (4) we assume that each map has been calibrated to have unit response to the lensed CMB. The lensed CMB in each frequency channel is thus determined by a single template T_{ℓ}^{CMB} . Similarly, we assume that each foreground s follows a single spatial template s_{ℓ} , up to a frequency dependence \mathbf{A}_s . Our normalized linear estimator for the lensed CMB takes the form $\hat{T}_{\ell} = \mathbf{w}_{\ell}^T \mathbf{T}_{\ell}$ for some weights \mathbf{w}_{ℓ} , which we are free to choose subject to the constraint $\mathbf{w}_{\ell}^T \mathbf{1} = 1$. The noise of the estimator \hat{T}_{ℓ} is

$$C_{\ell}^{\text{tot}} = \mathbf{w}_{\ell}^T \mathbf{C}_{\ell} \mathbf{w}_{\ell}, \quad (5)$$

where $(2\pi)^2 \delta_0^D \mathbf{C}_{\ell} = \langle \mathbf{T}_{\ell} \mathbf{T}_{\ell}^T \rangle$ is the covariance matrix of the maps \mathbf{T}_{ℓ} .

In this work we consider a Simons Observatory-like experiment, which will observe the CMB in six frequency channels: 27, 39, 93, 145, 225, and 280 GHz. We follow [17] in modeling the detector noise and atmospheric contributions to \mathbf{C}_{ℓ} , assuming the “goal” noise levels. The theory power spectra and SEDs for the galactic dust and extragalactic foregrounds are taken from [18].

Before deriving the optimal linear combination for CMB lensing, we start by reviewing the case of the standard ILC and deprojection. These two limiting cases, commonly used in CMB data analysis, correspond to two extremes. The standard ILC focuses only on the total map noise, without trying to reduce any specific foreground bias, whereas deprojection completely nulls a given foreground (in so far as its frequency dependence is known exactly), regardless of the noise penalty. As we show below, our CMB lensing-optimized ILC is a compromise between these two extremes.

A. Standard ILC

The standard harmonic ILC is defined to minimize the total map variance from noise and foregrounds. Minimizing C_{ℓ}^{tot} subject to the constraint $\mathbf{w}_{\ell}^T \mathbf{1} = 1$ can be achieved by introducing a Lagrange multiplier λ_{ℓ} and minimizing the Lagrangian:

$$\mathcal{L}[\mathbf{w}_{\ell}, \lambda_{\ell}] \equiv C_{\ell}^{\text{tot}} + \lambda_{\ell} (1 - \mathbf{w}_{\ell}^T \mathbf{1}). \quad (6)$$

Setting $\partial \mathcal{L} / \partial \mathbf{w}_{\ell} = 0$ forces $\mathbf{w}_{\ell} \propto \mathbf{C}_{\ell}^{-1} \mathbf{1}$. The normalization is fixed by the constraint $\mathbf{w}_{\ell}^T \mathbf{1} = 1$, from which we recover the standard result [19, 20]:

$$\mathbf{w}_{\ell}^{\text{ILC}} = \frac{\mathbf{C}_{\ell}^{-1} \mathbf{1}}{\mathbf{1}^T \mathbf{C}_{\ell}^{-1} \mathbf{1}}. \quad (7)$$

While the ILC minimizes the total power spectrum of the combined map, there is no guarantee that the power spectrum of individual foreground components will be reduced, so long that the foregrounds are a subdominant source of power. Even if the foreground power spectra are reduced, the ILC it is not designed to reduce the bispectrum or trispectrum of the foregrounds, which in turn source the biases to lensing. Thus it is possible for a foreground to have a subdominant power spectrum, but a dominant bispectrum or trispectrum, in which case the ILC does nothing to suppress the resulting bias.

B. Deprojection

In certain applications, one may need to null a specific foreground. For instance, tSZ measurements (y -maps and stacks) are contaminated by thermal dust emission, if unaccounted for. In this context, one may use constrained ILC maps with CIB deprojection [21, 22]. As shown in Fig. 1, one can deproject one or several foregrounds to remove the corresponding biases, at the cost of increasing the temperature map noise. Deprojection amounts to choosing the weights \mathbf{w}_{ℓ} to have zero response to foreground s : $\mathbf{w}_{\ell}^T \mathbf{A}_s = 0$. Minimizing the map noise subject to this additional constraint leads to minimizing

$$\mathcal{L}[\mathbf{w}_{\ell}, \lambda_{\ell}, \sigma_{\ell}] \equiv C_{\ell}^{\text{tot}} + \lambda_{\ell} (1 - \mathbf{w}_{\ell}^T \mathbf{1}) - \sigma_{\ell} \mathbf{w}_{\ell}^T \mathbf{A}_s. \quad (8)$$

Setting $\partial\mathcal{L}/\partial\mathbf{w}_\ell = 0$ forces $\mathbf{w}_\ell \propto \mathbf{C}_\ell^{-1}(\mathbf{1} + \sigma_\ell \mathbf{A}_s/\lambda_\ell)$. Enforcing the constraint $\mathbf{w}_\ell^T \mathbf{A}_s = 0$ fixes the coefficient:

$$\sigma_\ell/\lambda_\ell = -\frac{\mathbf{1}^T \mathbf{C}_\ell^{-1} \mathbf{A}_s}{\mathbf{A}_s^T \mathbf{C}_\ell^{-1} \mathbf{A}_s}. \quad (9)$$

Finally, $\mathbf{w}_\ell^T \mathbf{1} = 1$ fixes the normalization. This technique can be trivially extended to simultaneously deproject more than one foreground ($\mathbf{w}_\ell^T \mathbf{A}_{s_i} = 0$ for $i = 1, \dots, n$) by introducing additional Lagrange multipliers.

IV. RECONSTRUCTION NOISE

A general quadratic estimator $\hat{\kappa}_L$ of the lensing convergence κ_L takes the form

$$\hat{\kappa}_L[T, T'] = N_L \int \frac{d^2\ell}{(2\pi)^2} F_{\ell, L-\ell} T_\ell T'_{L-\ell} \quad (10)$$

where N_L is some normalization, chosen so that the estimator has unit response to the signal ($\langle \hat{\kappa}_L \rangle = \kappa_L + \text{bias}$, and $F_{\ell, L-\ell}$ are the weights, which are in principle arbitrary. In this work, both maps T and T' have unit response to the CMB, such that the normalization is

$$N_L^{-1} = \int \frac{d^2\ell}{(2\pi)^2} F_{\ell, L-\ell} f_{\ell, L-\ell}^\kappa. \quad (11)$$

This estimator applied to the map \hat{T}_ℓ has noise

$$\begin{aligned} \mathcal{N}_L[\mathbf{w}] &= \langle \hat{\kappa}_L[\hat{T}, \hat{T}] \hat{\kappa}_{-L}[\hat{T}, \hat{T}] \rangle / (2\pi)^2 \delta_0^D \\ &= 2N_L^2 \int \frac{d^2\ell}{(2\pi)^2} F_{\ell, L-\ell}^2 C_\ell^{\text{tot}} C_{|L-\ell|}^{\text{tot}} \end{aligned} \quad (12)$$

for $L > 0$. In going from the first to second line of Eq. (12), we have only kept the disconnected contributions to the auto-correlation, as is customary in lensing noise $N^{(0)}$ calculations. We also assumed that F is symmetric and even in both of its arguments.

Note that both the normalization N_L and the weights $F_{\ell, L-\ell}$ could in principle depend on the weights \mathbf{w}_ℓ . In this analysis, the spatial weights $F_{\ell, L-\ell}$ depend on the map noise C_ℓ^{tot} , which in turn depend on the frequency weights \mathbf{w}_ℓ . In App. B, we show that the ILC is the frequency combination which minimizes the lensing noise, for all the estimators we consider here. The top panel of Fig. 2 shows the lensing noise for the standard QE for various frequency combinations, with $\ell_{\text{max}, T} = 3000$.

A. Lensing amplitude noise

The lensing amplitude A_{lens} is defined to be the ratio of the measured convergence auto-spectrum to the truth, with fiducial value 1. An estimator for the lensing amplitude can be easily constructed for each mode L : $\hat{A}_{\text{lens}, L} = \hat{C}_L^\kappa / C_L^\kappa$, where \hat{C}_L^κ denotes the measured

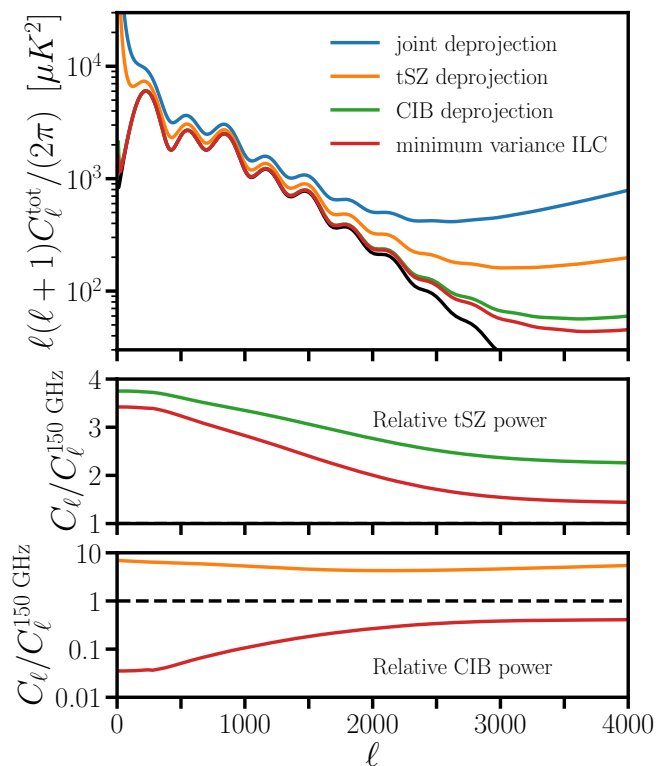


Figure 1. The top panel shows the temperature power spectrum for different multi-frequency combinations, compared to the noiseless lensed CMB (black curve). Compared to the minimum-variance combination (ILC, red), the deprojection of tSZ (orange) comes at a larger cost in noise than that of CIB (green). Jointly deprojecting tSZ and CIB (blue) causes a more than ten-fold noise increase at $\ell \sim 3000$, where most of the lensing information resides. We note that the significant increase in power at large scales from joint or tSZ deprojection is primarily due to a large increase in atmospheric noise. The middle and bottom panels show the tSZ and CIB power spectra, respectively, compared to the single-frequency case (150 GHz). For tSZ (middle), the standard ILC and the CIB deprojection cause a significant increase. For CIB (bottom), the ILC very effectively reduces the power; however, the tSZ-deprojection enhances it by more than an order of magnitude.

convergence auto-spectrum. Given a range of modes, the minimum variance estimator takes the form

$$\hat{A}_{\text{lens}} = \left(\int \frac{d^2\mathbf{L}}{(2\pi)^2} \frac{\hat{C}_L^\kappa (C_L^\kappa)^2}{C_L^\kappa \sigma_L^2} \right) / \left(\int \frac{d^2\mathbf{L}}{(2\pi)^2} \frac{(C_L^\kappa)^2}{\sigma_L^2} \right), \quad (13)$$

where $\sigma_L^2 = 2(C_L^\kappa + N_L^\kappa)^2$. This is simply an inverse variance weighting of the estimators for each Fourier mode.

Here we simply aim to minimize the lensing noise:

$$\sigma_{A_{\text{lens}}}^2 = \frac{1}{4\pi f_{\text{sky}}} \left(\int \frac{d^2\mathbf{L}}{(2\pi)^2} \frac{(C_L^\kappa)^2}{\sigma_L^2} \right)^{-1} \quad (14)$$

Minimizing $\sigma_{A_{\text{lens}}}$ is equivalent to maximizing the integral in Eq. (14), which has a strictly positive integrand. Thus maximizing the integral is equivalent to maximizing the integrand for each \mathbf{L} , which in turn is equivalent to minimizing σ_L^2 , or N_L^κ . Therefore the weights that minimize

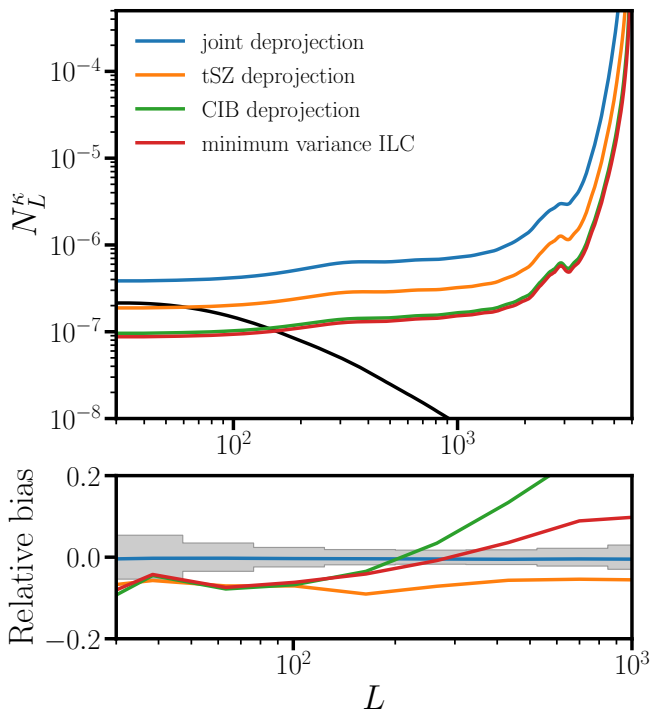


Figure 2. As illustrated in the top panel, the standard ILC is the linear combination of multi-frequency maps which minimizes the lensing noise, for the standard QE and all the other estimators considered in this paper. As expected from Fig. 1, deprojecting CIB (green) comes at almost no cost in lensing noise, whereas deprojecting tSZ (yellow) multiplies the noise by ~ 2 . Despite causing a more than ten-fold increase in temperature noise, joint deprojection only causes a four-fold increase in lensing noise. This can be understood by counting the number of signal-dominated temperature modes. The bottom panel shows the fractional bias to the CMB lensing auto-spectrum for the various frequency combinations. The gray boxes denote the bandpower errors for the standard ILC. For the ILC, the bias is largely dominated by the enhanced tSZ signal. While the joint deprojection successfully nulls all biases, simply deprojecting tSZ or CIB does not.

the noise of the lensing amplitude $\sigma_{A_{\text{lens}}}$ are equivalent to the weights that minimize the lensing reconstruction noise N_L^κ , which is given by the standard ILC.

V. BIAS TO LENSING RECONSTRUCTION FROM SIMULATIONS

A. Method

To estimate the non-Gaussian lensing biases from extragalactic foregrounds, we follow the procedure in [8, 10]. We ignore the lensed-foreground bias, due to the fact that the foregrounds are emitted at cosmological distances and are themselves lensed [23]. The non-Gaussian foreground bias to the lensing auto-spectrum can be decomposed into

a primary, secondary and trispectrum terms [4, 5, 7, 8]:

$$\begin{aligned} \langle \mathcal{Q}\mathcal{Q} \rangle = & \underbrace{\langle \kappa_{\text{CMB}} \kappa_{\text{CMB}} \rangle}_{\text{lensing signal}} + 2 \underbrace{\langle \mathcal{Q}[T^{\text{CMB}}, T^{\text{CMB}}] \mathcal{Q}[s, s] \rangle}_{\text{Primary bispectrum bias}} \\ & + 4 \underbrace{\langle \mathcal{Q}[T^{\text{CMB}}, s] \mathcal{Q}[T^{\text{CMB}}, s] \rangle}_{\text{Secondary bispectrum bias}} + \underbrace{\langle \mathcal{Q}[s, s] \mathcal{Q}[s, s] \rangle}_{\text{Trispectrum bias}}, \end{aligned} \quad (15)$$

where $\mathcal{Q}[T, T']$ represents any (symmetric) quadratic estimator evaluated on the maps T and T' . For the cross-correlation of CMB lensing with an external tracer g of the mass (e.g., galaxies or galaxy lensing), only the primary bispectrum bias appears:

$$\langle g\mathcal{Q} \rangle = \underbrace{\langle g\kappa_{\text{CMB}} \rangle}_{\text{lensing signal}} + \underbrace{\langle g\mathcal{Q}[s, s] \rangle}_{\text{bispectrum bias}}. \quad (16)$$

In both cases, the bispectrum biases (primary and secondary) appear because the foregrounds are simultaneously non-Gaussian and correlated with the true CMB lensing signal. The trispectrum bias in auto-correlation is present for any non-Gaussian foreground, regardless of any potential correlation with CMB lensing.

In what follows, we estimate each of these bias terms separately for the total contribution from extragalactic foregrounds (tSZ + CIB + kSZ + radio PS). We apply the method presented in [8] identically. This uses the non-Gaussian foreground simulations of [24] at 150 GHz, which are then rescaled to any other frequency, given the expected SED of each component from [18], as implemented in *MicroCoSM*². We choose the tracer g to be a LSST-like sample constructed from reweighting the simulated halos, as in [8, 10], to obtain the correct redshift distribution and linear bias.

In all cases, the foreground maps from [24] are masked before the lensing reconstruction. The mask is obtained by detecting the individual point sources with flux density higher than 5 mJy in the 150 GHz map using a matched filter, and masking the foreground map with a 3 arcmin radius disks around each source. We then add the unmasked lensed CMB map to the foreground map. This simulates the effect of inpainting, or source template subtraction, as expected for SO. We also explore a more aggressive masking strategy in App. D, which could reduce the bias to the profile hardened estimator by as much as 30%, provided that the effects of the mean field produced by this more aggressive mask are kept under control. Masking the point sources may also mask the high κ peaks, which would bias the signal low [25, 26], but inpainting or template subtraction should circumvent this problem.

Another important method for avoiding foreground contamination is to discard the highest multipoles in the input temperature map. Unless otherwise indicated, we assume that only modes with $\ell \leq \ell_{\text{max}, T} = 3000$ are included. In App. E, we explore the sensitivity of our results to this choice.

² <https://github.com/EmmanuelSchaan/MicroCoSM>

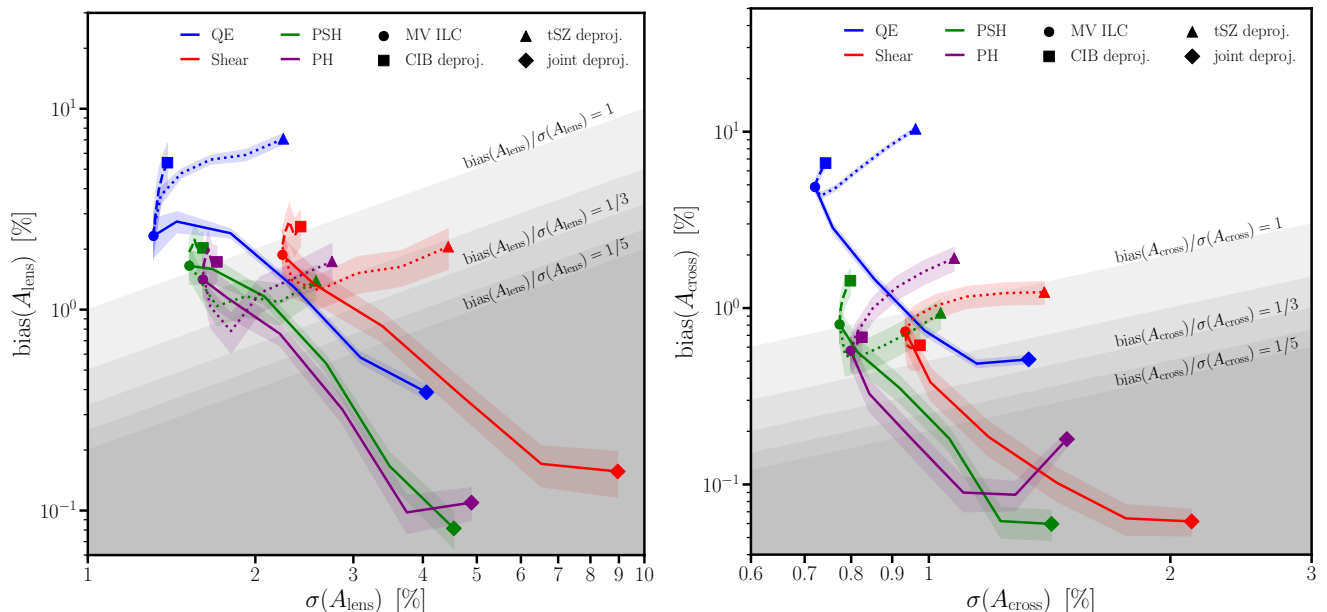


Figure 3. *Left:* The total |bias| and noise on the lensing amplitude A_{lens} for the various ILC combinations. The biases for the standard QE are shown in blue, Shear in red, the Point Source Hardened estimator in green, and the Profile Hardened estimator in purple. *Right:* The same as the left panel, but for the amplitude A_{cross} of the cross-correlation with the mock LSST sample. Circles denote the minimum variance (MV) ILC, squares (triangles) denote CIB (tSZ) deprojection, while diamonds denote the simultaneous deprojection of both the CIB and tSZ. The lines connecting MV ILC to either CIB, tSZ or joint deprojection follow Eq. (19), with step sizes $\Delta t = 0.2$. Throughout we sum over $20 < L < 1000$ when calculating the bias and noise for either A_{lens} or A_{cross} , and take $f_{\text{sky}} = 0.4$.

B. Lensing amplitude bias

The bias to the lensing amplitude A_{lens} is given by

$$\text{bias}(A_{\text{lens}}) = \frac{\int d^2 \mathbf{L} C_L^\kappa \text{bias}(C_L^\kappa) / \sigma_L^2}{\int d^2 \mathbf{L} (C_L^\kappa)^2 / \sigma_L^2}, \quad (17)$$

where the bias to the lensing auto-correlation is the sum of the primary, secondary and trispectrum terms, as defined in Eq. (15). Analogous expressions for the noise and bias of the amplitude A_{cross} of the cross-correlation of CMB lensing with an external tracer g may be obtained by substituting C_L^κ with $C_L^{\kappa g}$ and setting $\sigma_L^2 = (C_L^\kappa + N_\ell^\kappa)(C_L^g + N_\ell^g) + (C_L^{\kappa g})^2$ in Eqs. (14) and (17).

We note that cancellations could potentially occur in Eq. (17) between the biases at different lensing multipoles. In App. A we explore how our optimal ILC weights change if we replace $\text{bias}(C_L^\kappa)$ with $|\text{bias}(C_L^\kappa)|$ in the expression above, which removes these cancellations.

VI. MINIMIZING LENSING BIAS AND NOISE

In this section we piece together the preceding results and show that an optimal multi-frequency linear combination for minimizing a combination of noise and bias (for both A_{lens} and A_{cross}) lies between the standard ILC and joint deprojection.

A. Method

Ideally, one would solve for the scale and frequency-dependent weights \mathbf{w}_ℓ which minimize some function \mathcal{L} of the variance and bias of A_{lens} (or A_{cross}). A reasonable choice for \mathcal{L} is:

$$\mathcal{L}[\mathbf{w}_\ell] = \sigma^2(A_{\text{lens}}) + f_b^2 \text{bias}^2(A_{\text{lens}}), \quad (18)$$

where f_b is a tuning parameter for favoring either noise ($f_b = 0$) or bias ($f_b \rightarrow \infty$). In practice, however, solving for \mathbf{w}_ℓ via a brute force minimization of \mathcal{L} is computationally intractable. The problem can be greatly simplified with a few well-motivated approximations; however, these turn out not to be accurate enough for our purposes (see App. C). We circumvent this issue with the following ansatz:

$$\mathbf{w}_\ell(t) = t \mathbf{X}_\ell + (1-t) \mathbf{w}_\ell^{\text{ILC}}, \quad (19)$$

which smoothly connects the standard ILC to some form of deprojection ($\mathbf{X}_\ell = \mathbf{w}_\ell^{\text{CIB deproj.}}$, $\mathbf{w}_\ell^{\text{tSZ deproj.}}$, or $\mathbf{w}_\ell^{\text{joint deproj.}}$), and reduces the dimensionality of the parameter space from $5 \times (\# \text{ of } \ell\text{-bins})$ to 1. In what follows we compute the bias and noise for each estimator along the line segment $t \in [0, 1]$, in steps of $\Delta t = 0.2$. The optimal value of t could in principle be found through minimizing Eq. (18) for fixed f_b . However, the value of f_b is arbitrarily chosen to reduce the bias to some desirable level, which can just as well be accomplished by choosing t to be the smallest value such that $\text{bias}(A_{\text{lens}})/\sigma(A_{\text{lens}})$

is below some desirable cutoff. This is similar in spirit to the approach taken in deriving the partially constrained ILC [27]. However, our cutoff is implemented directly at the bias level, while the cutoff in [27] is implemented at the foreground power level.

B. Results

Our results are summarized in Fig. 3, where we plot the bias vs noise for each estimator as we vary the weights w_ℓ between the standard ILC and some form of deprojection \mathbf{X}_ℓ , following the prescription of Eq. (19). We first note that deprojecting a single foreground (tSZ or CIB) often boosts the overall bias. As shown in bottom two panels Fig. 1, deprojecting tSZ significantly boosts the CIB power spectrum (and vice versa), increasing the overall bias for both the auto and cross-correlation in many cases.

We find that hardening against a tSZ-like profile outperforms the remaining estimators for any bias/noise cutoff³, both in auto- and cross-correlation. Profile hardening at ILC ($t = 0$, purple circles in Fig. 3) reduces the bias to the lensing amplitude by 40%, at a 20% cost in noise, while the bias to the cross-correlation is reduced by nearly an order of magnitude at a 10% noise cost, relative to the standard quadratic estimator at ILC (blue circles in Fig. 3).

For all estimators, comfortably reducing the bias below $\sigma(A_{\text{lens}})/2$ (or $\sigma(A_{\text{cross}})/2$) requires partially deprojecting both the tSZ and CIB. We find that the noise cost for partial joint deprojection is much more significant for the auto-correlation than for the cross-correlation. In particular, for the profile hardened estimator, reducing the bias below $\sigma/2$ via partial joint deprojection corresponds to $t \sim 0.2$, resulting in a 20% and 5% noise penalty for the auto- and cross-correlation respectively, relative to profile hardening at ILC⁴. Further improvements can be made by pushing to a higher $\ell_{\text{max},T}$ and moving closer to deprojection (higher t), or through a more aggressive masking scheme, as discussed in Apps. E and D respectively.

VII. CONCLUSIONS

CMB lensing is quickly becoming one of the most powerful cosmological tools. Percent level measurements from the next generation of experiments require excellent control of extragalactic foregrounds, especially from temperature anisotropies. Traditional analyses have used

multi-frequency foreground reduction methods when possible. Here we combined geometric methods with multi-frequency information to optimize the lensing reconstruction. In doing so, we maximize the amount of information that can be extracted and the robustness of the measurement, while minimizing the potential biases. We also note that any multi-frequency linear combination alone will not reduce the bias from kSZ, since it preserves the black-body nature of the CMB. While kSZ bias can be significant [7], it is quite effectively reduced by the shear-only or bias hardened estimators [8, 10]. Moreover, direct measurement of kSZ by combining low-redshift galaxy catalogs together with high-resolution CMB maps [28] can potentially provide templates for subtraction at the map level. We find that the standard ILC can lead to significant biases, which can be partially mitigated with geometric methods such as profile hardening, which alone can reduce the bias by 40% (in auto-correlation) or almost an order of magnitude (in cross-correlation) with minimal (10 – 20%) increase in noise. The bias can be further reduced to any desirable bias/noise cutoff via partial joint deprojection of the tSZ and CIB, with the auto-correlation suffering a higher relative noise penalty (factor of ~ 4) than the cross-correlation. We find that deprojection of tSZ or CIB alone do not reduce the biases, due to an enhancement in the component that is not deprojected. Finally, we have explored possible improvements with more aggressive masking (App. D) and varying $\ell_{\text{max},T}$ (App. E).

While here we have focused on temperature reconstruction, similar bias hardening techniques should be applicable to extragalactic foregrounds in polarization as well, where the biases are expected to be less severe. We leave a full exploration of this to future work.

ACKNOWLEDGMENTS

We thank Anthony Challinor, Jo Dunkley, Giulio Fabbian, Colin Hill, Akito Kusaka, Antony Lewis, Mathew Madhavacheril, Neelima Sehgal, Alexander Van Engelen and Martin White for useful discussions in the preparation of this work. N.S. is supported by the NSF. E.S. is supported by the Chamberlain fellowship at Lawrence Berkeley National Laboratory. S.F. is supported by the Physics Division of Lawrence Berkeley National Laboratory. This work used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

³ This statement is no longer true near joint deprojection ($t \simeq 1$). In this regime, the PSH estimator has a lower bias (and noise) than the PH estimator, since in this case the foreground biases are primarily sourced by radio point sources (as opposed to tSZ or CIB), which the former estimator is designed to null. In practice, however, one does not need to push to $t \simeq 1$ to reduce the bias

to an acceptable level.

⁴ Relative to the standard QE at ILC, the noise penalty for reducing the bias below $\sigma/2$ using profile hardening and partial joint deprojection is 50% and 15% for the auto- and cross-correlation respectively.

- [1] A. Blanchard and J. Schneider, *Astron. Astrophys.* **184**, 1 (1987).
- [2] A. Lewis and A. Challinor, *Phys. Rep.* **429**, 1 (2006).
- [3] D. Hanson, A. Challinor, and A. Lewis, *General Relativity and Gravitation* **42**, 2197 (2010).
- [4] A. van Engelen *et al.*, *Astrophys. J.* **786**, 13 (2014).
- [5] S. J. Osborne, D. Hanson, and O. Doré, *J. Cosm. Astrop. Phys.* **2014**, 024 (2014).
- [6] M. S. Madhavacheril and J. C. Hill, *Phys. Rev. D* **98**, 023534 (2018).
- [7] S. Ferraro and J. C. Hill, *Phys. Rev. D* **97**, 023512 (2018).
- [8] E. Schaan and S. Ferraro, *Phys. Rev. Lett.* **122**, 181301 (2019).
- [9] E. J. Baxter *et al.*, *Phys. Rev. D* **99**, 023508 (2019).
- [10] N. Sailer, E. Schaan, and S. Ferraro, *Phys. Rev. D* **102**, 063517 (2020).
- [11] W. Hu and T. Okamoto, *Astrophys. J.* **574**, 566 (2002).
- [12] J. Delabrouille and J. F. Cardoso, in *Data Analysis in Cosmology*, edited by V. J. Martínez, E. Saar, E. Martínez-González, and M. J. Pons-Bordería (PUBLISHER, ADDRESS, 2009), Vol. 665, pp. 159–205.
- [13] O. Darwish *et al.*, arXiv e-prints arXiv:2004.01139 (2020).
- [14] S. Patil, S. Raghunathan, and C. L. Reichardt, (2019).
- [15] O. Darwish *et al.*, arXiv e-prints arXiv:2111.00462 (2021).
- [16] T. Namikawa, D. Hanson, and R. Takahashi, *Mon. Not. R. Astron. Soc.* **431**, 609 (2013).
- [17] P. Ade *et al.*, *J. Cosm. Astrop. Phys.* **2019**, 056 (2019).
- [18] J. Dunkley *et al.*, *J. Cosm. Astrop. Phys.* **2013**, 025 (2013).
- [19] H. K. Eriksen, A. J. Banday, K. M. Gorski, and P. B. Lilje, *The Astrophysical Journal* **612**, 633–646 (2004).
- [20] M. Tegmark, A. de Oliveira-Costa, and A. J. S. Hamilton, *Physical Review D* **68**, (2003).
- [21] M. Remazeilles, J. Delabrouille, and J.-F. Cardoso, *Monthly Notices of the Royal Astronomical Society* **410**, 2481–2487 (2010).
- [22] M. S. Madhavacheril *et al.*, *Physical Review D* **102**, (2020).
- [23] N. Mishra and E. Schaan, *Phys. Rev. D* **100**, 123504 (2019).
- [24] N. Sehgal *et al.*, *Astrophys. J.* **709**, 920 (2010).
- [25] G. Fabbian, J. Carron, A. Lewis, and M. Lembo, *Physical Review D* **103**, (2021).
- [26] J. Liu and J. C. Hill, *Physical Review D* **92**, (2015).
- [27] Y. S. Abylkairov, O. Darwish, J. C. Hill, and B. D. Sherwin, *Physical Review D* **103**, (2021).
- [28] E. Schaan *et al.*, *Phys. Rev. D* **103**, 063513 (2021).
- [29] D. S. Y. Mak, A. Challinor, G. Efstathiou, and G. Lagache, *Mon. Not. R. Astron. Soc.* **466**, 286 (2017).
- [30] Planck Collaboration *et al.*, *Astron. Astrophys.* **536**, A18 (2011).
- [31] M. Lembo, G. Fabbian, J. Carron, and A. Lewis, arXiv e-prints arXiv:2109.13911 (2021).

Appendix A: Cancellations

As mentioned in §V, the definition of $\text{bias}(A_{\text{lens}})$ permits undesirable cancellations in biases across lensing multipoles. A simple method for avoiding these cancellations is to replace $\text{bias}(C_L^{\kappa})$ with $|\text{bias}(C_L^{\kappa})|$ in Eq. (17). Fig. 4 shows how our results change with this substitution. In particular for the profile hardened estimator, we find that one has to deproject out to $t \sim 0.4$ (as opposed to $t \sim 0.2$) in order to achieve $\text{bias}/\text{noise} < 1/2$.

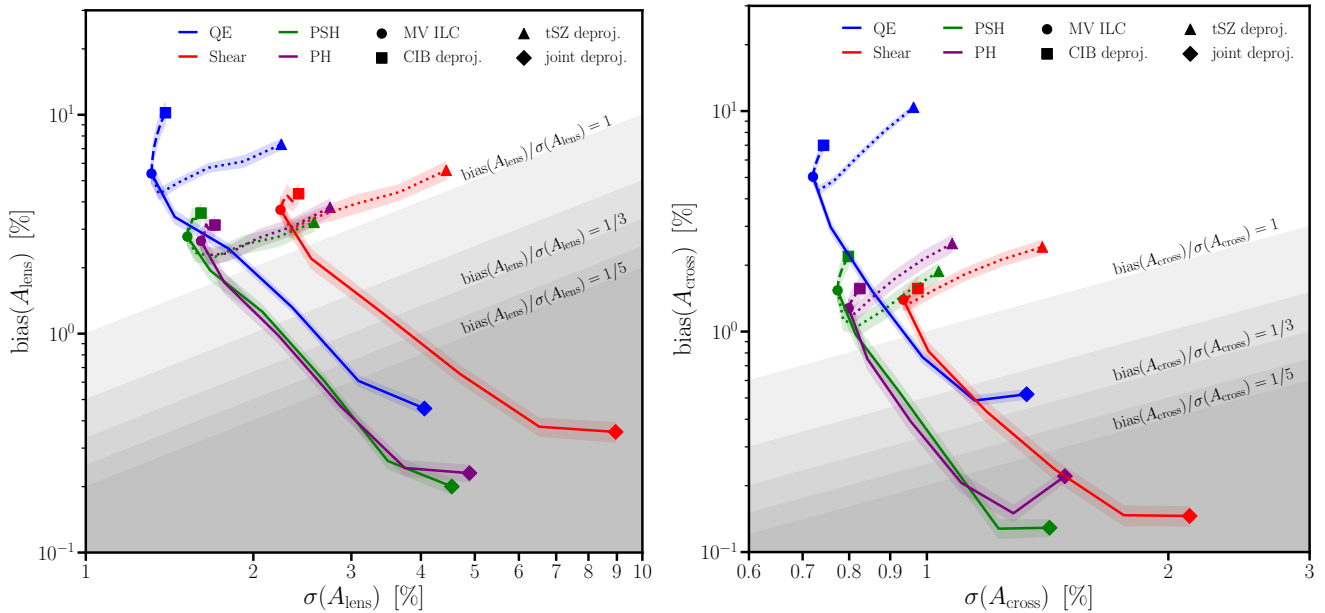


Figure 4. The same as Fig. 3, but with $|\text{bias}(C_L^{\kappa})|$ replacing $\text{bias}(C_L^{\kappa})$ in Eq. (17), and similarly for the cross-correlation with the mock LSST sample.

In addition to the aforementioned cancellations, for any given foreground, the primary, secondary and trispectrum terms may also cancel separately due to the linear combination of frequencies. While the distinction of primary, secondary and trispectrum terms is somewhat arbitrary, this cancellation relies on the relative sizes of the foreground bispectrum and trispectrum to be correct in the simulation used. Additionally, at any lensing multipole, the lensing biases from two distinct foregrounds may have different signs and partially cancel. While one can in principle take advantage of these cancellations, doing so relies on the simulations having the correct relative amplitude of each foreground, including their correlations. We will explore the sensitivity of our optimized ILC to these results in a future work, which will more realistically include polarization, in addition to temperature data.

Appendix B: Proof that ILC minimizes the lensing noise for all QEs

In this Appendix, we show that the linear combination of frequency maps which minimizes the lensing noise is the standard ILC, for all the estimators we consider. Specifically, we show the following: if the spatial weights $F_{\ell, \mathbf{L}-\ell}^\kappa$ of the estimator $\hat{\kappa}$ depend on the frequency weights \mathbf{w}_ℓ only through the map noise C_ℓ^{tot} , then the frequency weights \mathbf{w}_ℓ that extremize $\mathcal{N}_L[\mathbf{w}]$ are the ILC weights. To minimize the noise subject to the constraint $\mathbf{w}_\ell^T \mathbf{1} = 1$ we set

$$\begin{aligned} \frac{\delta \mathcal{N}_L[\mathbf{w}]}{\delta \mathbf{w}_\ell} &= 4N_L^2 F_{\ell, \mathbf{L}-\ell}^2 \left[C_\ell^{\text{tot}} + C_{|\mathbf{L}-\ell|}^{\text{tot}} \right] \mathbf{C}_\ell \mathbf{w}_\ell + 4N_L^2 \int \frac{d^2 \ell'}{(2\pi)^2} F_{\ell', \mathbf{L}-\ell'} \frac{\delta F_{\ell', \mathbf{L}-\ell'}}{\delta \mathbf{w}_\ell} C_{\ell'}^{\text{tot}} C_{|\mathbf{L}-\ell'|}^{\text{tot}} \\ &\quad - 2N_L \mathcal{N}_L[\mathbf{w}] \int \frac{d^2 \ell'}{(2\pi)^2} \frac{\delta F_{\ell', \mathbf{L}-\ell'}}{\delta \mathbf{w}_\ell} f_{\ell', \mathbf{L}-\ell'}^\kappa \\ &= \lambda_{\ell, \mathbf{L}} \mathbf{1}, \end{aligned} \quad (\text{B1})$$

where $\lambda_{\ell, \mathbf{L}}$ is some Lagrange multiplier. If the spatial weights $F_{\ell, \mathbf{L}-\ell}^\kappa$ only depend on frequency weights \mathbf{w}_ℓ through C_ℓ^{tot} and $C_{|\mathbf{L}-\ell|}^{\text{tot}}$, then

$$\begin{aligned} \frac{\delta F_{\ell', \mathbf{L}-\ell'}^\kappa}{\delta \mathbf{w}_\ell} &= \frac{\partial F_{\ell', \mathbf{L}-\ell'}^\kappa}{\partial C_{\ell'}^{\text{tot}}} \frac{\delta C_{\ell'}^{\text{tot}}}{\delta \mathbf{w}_\ell} + \frac{\partial F_{\ell', \mathbf{L}-\ell'}^\kappa}{\partial C_{|\mathbf{L}-\ell'|}^{\text{tot}}} \frac{\delta C_{|\mathbf{L}-\ell'|}^{\text{tot}}}{\delta \mathbf{w}_\ell} \\ &= (2\pi)^2 \frac{\partial F_{\ell, \mathbf{L}-\ell}^\kappa}{\partial C_\ell^{\text{tot}}} \left[\delta_{\ell-\ell'}^D + \delta_{\ell-\mathbf{L}+\ell'}^D \right] \mathbf{C}_\ell \mathbf{w}_\ell \\ &\equiv \mathcal{A}_{\ell, \ell', \mathbf{L}} \mathbf{C}_\ell \mathbf{w}_\ell \end{aligned} \quad (\text{B2})$$

where $\mathcal{A}_{\ell, \ell', \mathbf{L}}$ is some scalar function. Plugging this in to our equation for $\delta \mathcal{N}_L[\mathbf{w}]/\delta \mathbf{w}_\ell$ gives

$$\frac{\delta \mathcal{N}_L[\mathbf{w}]}{\delta \mathbf{w}_\ell} = \mathcal{B}_{\ell, \mathbf{L}} \mathbf{C}_\ell \mathbf{w}_\ell = \lambda_{\ell, \mathbf{L}} \mathbf{1} \quad (\text{B3})$$

for some scalar function $\mathcal{B}_{\ell, \mathbf{L}}$. Solving this equation for \mathbf{w}_ℓ gives $\mathbf{w}_\ell = \lambda_{\ell, \mathbf{L}} C_\ell^{-1} \mathbf{1} / \mathcal{B}_{\ell, \mathbf{L}}$. The Lagrange multiplier is chosen so the weights sum to unity, which yields the standard ILC weights, given by Eq. (7).

Note that the standard QE, bias-hardened and shear estimators have respective weights:

$$F_{\ell, \mathbf{L}-\ell}^\kappa = \frac{f_{\ell, \mathbf{L}-\ell}^\kappa}{2C_\ell^{\text{tot}} C_{|\mathbf{L}-\ell|}^{\text{tot}}} \quad , \quad F_{\ell, \mathbf{L}-\ell}^{\text{BH}} = \frac{f_{\ell, \mathbf{L}-\ell}^\kappa - N_L^\kappa \mathcal{R}_L f_{\ell, \mathbf{L}-\ell}^s}{2C_\ell^{\text{tot}} C_{|\mathbf{L}-\ell|}^{\text{tot}}} \quad , \quad F_{\ell, \mathbf{L}-\ell}^{\text{shear}} = \cos(2\theta_{\mathbf{L}, \ell-\mathbf{L}/2}) \frac{C_\ell^0}{2(C_\ell^{\text{tot}})^2} \frac{d \ln C_\ell^0}{d \ln \ell} \quad , \quad (\text{B4})$$

which only depend on \mathbf{w}_ℓ through C_ℓ^{tot} . Thus their noises are minimized by the ILC weights.

Appendix C: Attempts at numerically minimizing the loss functions: approximations

Finding the temperature weights which minimize the desired combination of lensing noise and bias is a difficult problem. Indeed, the space of parameters to optimize is large, since the temperature weights are free functions of scale and frequency. For each choice of temperature weights, one needs to perform multiple lensing reconstructions on the foreground simulations to estimate the primary, secondary and trispectrum biases. This step is computationally expensive. This needs to be repeated for each choice of estimator (QE, shear, PSH, PH), and at each step when exploring the space of temperature weights. Below, we describe approximations which dramatically simplify the problem, and point out their inaccuracies.

One may assume that each foreground is determined by a single spatial template, whose amplitude is rescaled at each frequency (Approximation 1). This approximation is exact for the CMB and for the tSZ, ignoring relativistic corrections. It is also excellent for the CIB, whose maps at different frequencies are highly correlated [29, 30].

To simplify further, we note that most of the CMB lensing signal to noise comes from temperature multipoles around $\ell = 3000$. As a result, most of the foreground biases also comes from these multipoles. At these multipoles, the ILC weights are slowly varying. We may therefore neglect the scale dependence of the ILC weights \mathbf{w}^{ILC} , and look for scale-independent optimal weights \mathbf{w} (Approximation 2). This approximation dramatically reduces the dimensionality of the minimization problem, making it much more tractable. In this approximation, the primary $\mathcal{P}_{0,L}^s$, secondary $\mathcal{S}_{0,L}^s$ and trispectrum $\mathcal{T}_{0,L}^s$ biases to A_{lens} only need to be computed at one frequency (e.g., 150 GHz). They can then simply be rescaled to any other frequency. Furthermore, because the temperature weights are scale-independent, the bispectrum and trispectrum biases only need to be computed from simulations once. They can then be rescaled for any new choice of temperature weights. In short, approximations 1 and 2 allow us to compute the foreground biases from simulation only once, rather than having to do it for every frequency and for every choice of weights \mathbf{w}_ℓ .

Finally, we expect the optimal multi-frequency combination will be close to the standard ILC. We indeed show this to be true for the profile hardened estimator when exploring part of the parameter space, going from ILC to some form of deprojection. This motivates us to Taylor expand the noise about the ILC weights to second order in ϵ (Approximation 3): $\mathbf{w} = \mathbf{w}_{\ell=3000}^{\text{ILC}} + \epsilon$. This approximation dramatically speeds up the noise computation during the optimization step, and accurately predicts the variance and bias of A_{lens} except at 27 GHz. A simple method for getting around this is to set the 27 GHz weight to zero, in which case the Taylor expansion agrees with the true noise to the subpercent level for the ϵ 's relevant for achieving bias/noise $\sim 1/2$.

In this simplified case, the scaling of the foreground biases with the frequency weights is simply:

$$\text{bias}(C_L^s)[\mathbf{w}] = \sum_s [\mathcal{P}_{0,L}^s + \mathcal{S}_{0,L}^s] \left(\frac{\mathbf{A}_s \cdot \mathbf{w}}{A_{s,150}} \right)^2 + \sum_s \mathcal{T}_{0,L}^s \left(\frac{\mathbf{A}_s \cdot \mathbf{w}}{A_{s,150}} \right)^4 \quad (\text{C1})$$

where $A_{s,150}$ is the scale-factor for the 150 GHz channel. In practice, this expression can be easily generalized to include the lensing bias from correlated foregrounds, e.g., CIB and tSZ. In summary, we are thus looking for a handful of parameters $(\epsilon_\nu)_{\nu=1\dots 5}$, which minimize the sum of squared variance and squared bias on A_{lens} .

However, we found approximations 2 and 3 not to be sufficiently accurate. Assuming the weights to be scale-independent leads to a significant misestimation of the variance and bias of A_{lens} , at the $\mathcal{O}(1)$ level.

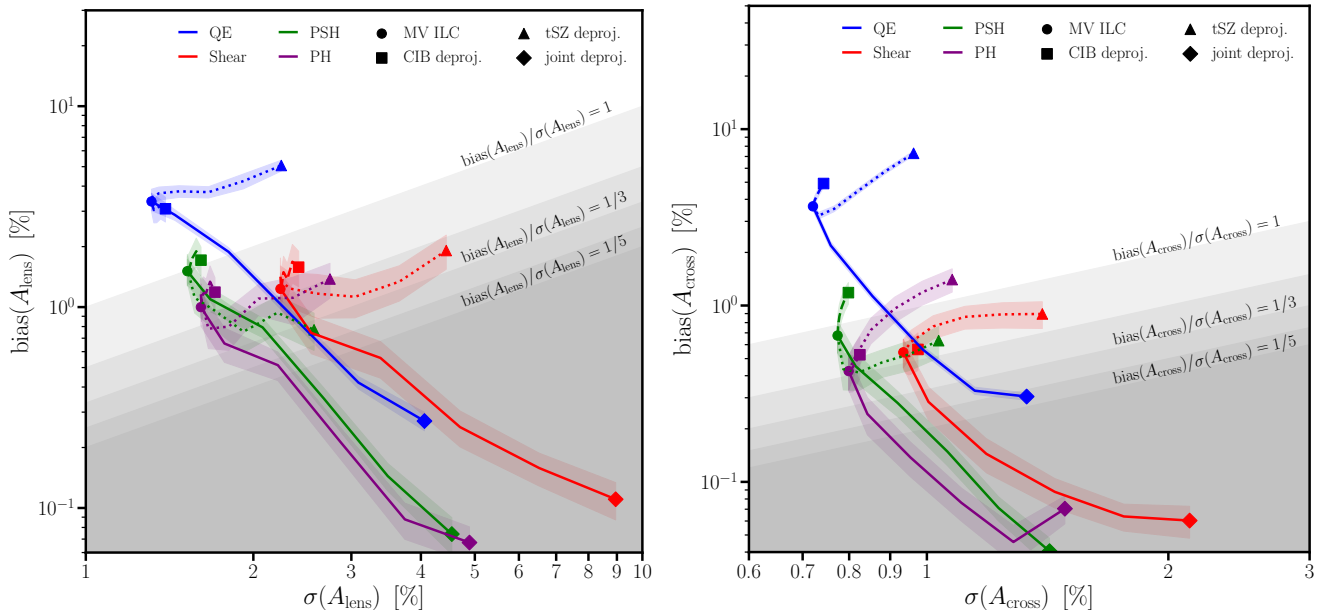


Figure 5. The same as Fig. 3, using the more aggressive masking technique described in App. D. When calculating σ , we neglect any differences in sky coverage due to more of the sky being masked (8% of the sky as opposed to 3% with the fiducial mask used in the main text), as this is a negligible effect.

Appendix D: Dependence of the lensing biases on the masking

We explore a more aggressive masking approach by running the matched filter at all the SO frequencies, masking only the sources (using disks with a 2 arcmin radius) that are detected with 5σ significance, and using the union of the masks at each frequency. This allows us to find much more IR (radio) sources, which are brighter at higher (lower)

frequencies. In an analysis on real data, this may make it more delicate to estimate the normalization and mean field of the lensing estimators [31], which has not been taken into account in our lensing biases calculations.

We note that in practice, masking changes the observed map covariance \mathbf{C}_ℓ , which in turn will influence the form of the ILC weights. For simplicity we have ignored this effect, and thus use the same ILC weights for the more aggressive masking approach as we did in the main text, which are derived using the theory foreground spectra from [18].

The resulting biases from this masking technique are shown in Fig. 5. In cross-correlation, the aggressive masking reduces the foreground bias for all estimators, as expected. In particular, at ILC ($t = 0$) the biases to the cross-correlation for PH, Shear, PSH, and QE are reduced by $\sim 30\%$, 40% , 10% , and 60% respectively, relative to the biases calculated using the single mask at 150 GHz. In auto-correlation, the biases for PH, Shear and PSH are similarly reduced using the more aggressive masking technique. However, the bias to the QE at ILC increases by 50% . This can be understood as follows. In auto-correlation, both the primary and trispectrum terms contribute to the foreground bias. The aggressive masking reduces both, but reduces the trispectrum much more. For the QE, a “lucky” cancellation allows the dominant negative primary bias (at low lensing- L) to cancel the dominant positive trispectrum bias (at high lensing- L). By reducing the trispectrum term more than the primary term, the aggressive masking spoils this cancellation, thus enhancing the QE bias in auto-correlation.

Appendix E: Dependence of the lensing biases on $\ell_{\max,T}$

In this section we explore how our results change with $\ell_{\max,T}$. As shown in Fig. 6, we still find that profile hardening outperforms the remaining estimators for any bias/noise cutoff and for any $\ell_{\max,T}$. We also find that small improvements are possible by pushing to a higher $\ell_{\max,T}$ with more joint deprojection. For the auto-correlation, the bias to the profile hardened estimator can be reduced below half of the statistical uncertainty at a 30% noise penalty by choosing $\ell_{\max,T} = 4000$ and $t \sim 0.4$ (as opposed to a 50% penalty with $\ell_{\max,T} = 3000$ and $t \sim 0.2$), relative to the standard QE at ILC with $\ell_{\max,T} = 3000$. For the cross-correlation, the noise penalty for achieving bias/noise $< 1/2$ with the profile hardened estimator can be reduced from 15% to 10% by similarly increasing $\ell_{\max,T}$ and t .

We note that for sufficiently high t , the biases and noises become roughly insensitive to the value of $\ell_{\max,T}$. This is due to the map noise blowing up at high- ℓ as the weights are pushed to joint deprojection, as shown in Fig. 1. This in turn down-weights high- ℓ modes and acts as an effective cutoff with $\ell_{\max}^{\text{eff}} < \ell_{\max,T}$. Instead of partial joint deprojection (Eq. (19)), one could naively use the standard ILC weights and reduce $\ell_{\max,T}$ to remove these high- ℓ modes, and hence the bias, in a potentially a more optimal way. However, in practice we find it more optimal to use a high $\ell_{\max,T}$ with joint deprojection than a lower $\ell_{\max,T}$ with the standard ILC weights. This is illustrated in Fig. 6, where we plot the bias and noise at ILC for $\ell_{\max,T} = 1500, 2000, \text{ and } 2500$ (individual circles). While lowering the $\ell_{\max,T}$ decreases the bias (as expected), the noise cost for doing so is much larger than increasing t with a higher $\ell_{\max,T}$ (as illustrated by the three lines for each estimator).

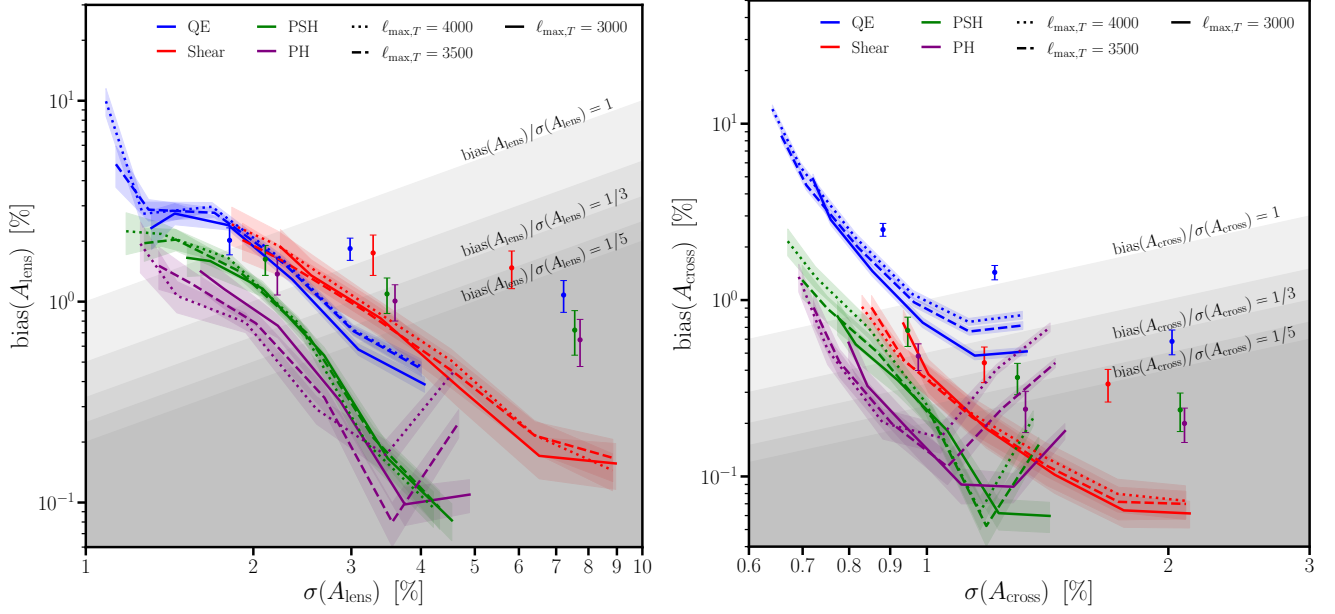


Figure 6. Noise vs bias for several $\ell_{\text{max},T}$'s. For $\ell_{\text{max},T} = 3000, 3500$, and 4000 we plot the line going from the standard ILC to joint deprojection, with the linestyle denoting the $\ell_{\text{max},T}$. For $\ell_{\text{max},T} = 1500, 2000$, and 2500 we plot the noise and bias at the standard ILC for each estimator (denoted by the circles). Note that for $\ell_{\text{max},T} = 1500$, the Shear estimator's noise is larger than 10% (3%) for the auto- (cross-) correlation, extending beyond the domain of the plot.