

UC Berkeley

UC Berkeley Previously Published Works

Title

Lower bias, lower noise CMB lensing with foreground-hardened estimators

Permalink

<https://escholarship.org/uc/item/5hw551rf>

Journal

Physical Review D, 102(6)

ISSN

2470-0010

Authors

Sailer, Noah

Schaan, Emmanuel

Ferraro, Simone

Publication Date

2020-09-15

DOI

10.1103/physrevd.102.063517

Peer reviewed

Lower bias, lower noise CMB lensing with foreground-hardened estimators

Noah Sailer,^{1,*} Emmanuel Schaen,^{2,1,†} and Simone Ferraro^{2,1,‡}

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*

Extragalactic foregrounds in temperature maps of the Cosmic Microwave Background (CMB) severely limit the ability of standard estimators to reconstruct the weak lensing potential. These foregrounds are not fully removable by multi-frequency cleaning or masking and can lead to large biases if not properly accounted for. For foregrounds made of a number of unclustered point sources, an estimator for the source amplitude can be derived and deprojected, removing any bias to the lensing reconstruction. We show with simulations that all of the extragalactic foregrounds in temperature can be approximated by a collection of sources with identical profiles, and that a simple bias hardening technique is effective at reducing any bias to lensing, at a minimal noise cost. We compare the performance and bias to other methods such as “shear-only” reconstruction, and discuss how to jointly deproject any arbitrary number of foregrounds, each with an arbitrary profile. In particular, for a Simons Observatory-like experiment foreground-hardened estimators allow us to extend the maximum multipole used in the reconstruction, increasing the overall statistical power by $\sim 50\%$ over the standard quadratic estimator, both in auto and cross-correlation. We conclude that source hardening outperforms the standard lensing quadratic estimator both in auto and cross-correlation, and in terms of lensing signal-to-noise and foreground bias.

I. INTRODUCTION

CMB lensing [1, 2] is rapidly becoming one of the most powerful tools in a cosmologist’s toolkit. With the ability to measure matter fluctuations to high redshift, no photometric redshift uncertainties and its sensitivity to crucial cosmological parameters, sub-percent precision measurements have the potential to greatly increase our understanding of the Universe.

Current and upcoming wide-field CMB experiments such as AdvACT [3], SPT-3G [4] and Simons Observatory [5] will heavily rely on reconstruction from CMB temperature fluctuations (as opposed to polarization). For future polarization-dominated experiments like CMB-S4 [6] and CMB-HD [7, 8], temperature lensing will still contribute a non-negligible part of the total signal-to-noise.

Temperature maps contain significant contamination from extragalactic foregrounds such as the thermal and kinematic Sunyaev-Zel’dovich effects (tSZ and kSZ), the Cosmic Infrared Background (CIB), and both radio and infrared point sources. Lensing reconstruction on contaminated maps creates significant biases in both the auto- and cross-correlation of the reconstructed field, which in turn can bias our inference of cosmology [9–14]. These biases can be up to several tens of percent for upcoming surveys, thus severely limiting our ability to use CMB lensing for precision cosmology.

Several techniques have been proposed and implemented to mitigate the impact of extragalactic foregrounds. Some are based on symmetries, such as the recently proposed shear-only reconstruction [13] (and related generalizations such as hybrid and general multipole estimators), while others make use of multi-frequency

maps to reduce the impact of particular foregrounds [11, 15].

Here we revisit and extend a bias hardening technique first proposed in [10, 16] and applied to real data in [17], and compare it to shear-only reconstruction and the hybrid estimator of [13]. We show that if we approximate extragalactic foregrounds as a collection of sources with identical profiles, an estimator for their amplitude can be obtained and hardened against in an optimal way to mitigate their impact on the reconstructed lensing convergence. With the use of realistic and correlated simulations of all of the foregrounds, we show that this approximation is excellent and leads to a very large suppression in bias, allowing for the use of considerably smaller scales in the reconstruction.

The remainder of the paper is organized as follows: In Section II we review the general formalism for bias hardening against a single foreground, then generalize the procedure to an arbitrary number of foregrounds, each with arbitrary profiles. In Section III we specialize to the case where foregrounds are point sources, and in Section IV we numerically explore the biases and the performance of mitigation through bias hardening or shear-only estimators. We conclude with Section V.

II. FOREGROUND-HARDENED QUADRATIC ESTIMATORS

In this section, we first review the derivation of the standard lensing quadratic estimator (QE), and rephrase it in terms of the linear response of the map covariance to the lensing convergence. We build on this formalism to construct foreground estimators, and finally null the linear response of the lensing QE to these foregrounds, as in [10, 16].

* nsailer@berkeley.edu

† eschaan@lbl.gov

‡ sferraro@lbl.gov

A. Lensing quadratic estimator, bispectrum & linear response

The observed CMB temperature T is the sum of the *lensed* primary CMB T^{CMB} and the foregrounds s , which will be approximated by a collection of sources in this work: $T = T^{\text{CMB}} + s$. The lensed CMB temperature T^{CMB} field is statistically invariant under translations. As a result, the off-diagonal covariance $\langle T_{\ell}^{\text{CMB}} T_{\mathbf{L}-\ell}^{\text{CMB}} \rangle$ with $\mathbf{L} \neq 0$ is zero. However, fixing the lensing convergence Fourier mode $\kappa_{\mathbf{L}}$ breaks this statistical homogeneity, and produces non-zero off-diagonal covariances:

$$\begin{aligned} \langle T_{\ell}^{\text{CMB}} T_{\mathbf{L}-\ell}^{\text{CMB}} \rangle_{\text{at fixed } \kappa_{\mathbf{L}}} &= f_{\ell, \mathbf{L}-\ell}^{\kappa} \kappa_{\mathbf{L}} + \mathcal{O}(\kappa_{\mathbf{L}}^2) \\ \text{with } f_{\ell, \mathbf{L}-\ell}^{\kappa} &\equiv \frac{2\mathbf{L}}{L^2} \cdot \left[\ell C_{\ell}^0 + (\mathbf{L} - \ell) C_{|\mathbf{L}-\ell|}^0 \right], \end{aligned} \quad (1)$$

where C_{ℓ}^0 is the *unlensed* power spectrum. From this, we obtain an unbiased lensing estimator (to first order in κ) from the ratio $T_{\ell} T_{\mathbf{L}-\ell} / f_{\ell, \mathbf{L}-\ell}^{\kappa}$. The standard lensing QE [18] is simply the inverse-variance weighted average of all these ‘‘building block’’ estimators, summing over ℓ at fixed \mathbf{L} . This leads to the usual result:

$$\begin{aligned} \hat{\kappa}_{\mathbf{L}} &= N_{\mathbf{L}}^{\kappa} \int \frac{d^2\ell}{(2\pi)^2} F_{\ell, \mathbf{L}-\ell}^{\kappa} T_{\ell} T_{\mathbf{L}-\ell} \\ (N_{\mathbf{L}}^{\kappa})^{-1} &= \int \frac{d^2\ell}{(2\pi)^2} F_{\ell, \mathbf{L}-\ell}^{\kappa} f_{\ell, \mathbf{L}-\ell}^{\kappa} \\ F_{\ell, \mathbf{L}-\ell}^{\kappa} &= \frac{f_{\ell, \mathbf{L}-\ell}^{\kappa}}{2C_{\ell}^{\text{tot}} C_{|\mathbf{L}-\ell|}^{\text{tot}}}. \end{aligned} \quad (2)$$

In order to generalize this construction to foreground quadratic estimators, we consider Eq. (1) and note that $f_{\ell, \mathbf{L}-\ell}^{\kappa}$ can be seen as the linear response of the two point correlator $\langle T_{\ell} T_{\mathbf{L}-\ell} \rangle_{\text{at fixed } \kappa_{\mathbf{L}}}$ to the lensing convergence mode $\kappa_{\mathbf{L}}$:

$$f_{\ell, \mathbf{L}-\ell}^{\kappa} = \frac{\delta \langle T_{\ell}^{\text{CMB}} T_{\mathbf{L}-\ell}^{\text{CMB}} \rangle_{\text{at fixed } \kappa_{\mathbf{L}}}}{\delta \kappa_{\mathbf{L}}} (\kappa_{\mathbf{L}} = 0). \quad (3)$$

Furthermore, by multiplying both sides of Eq. (1) with $\kappa_{-\mathbf{L}}$, we see that this linear response can be expressed as the ratio of the following bispectrum and power spectrum:

$$f_{\ell, \mathbf{L}-\ell}^{\kappa} = \frac{\langle T_{\ell}^{\text{CMB}} T_{\mathbf{L}-\ell}^{\text{CMB}} \kappa_{-\mathbf{L}} \rangle}{\langle \kappa_{\mathbf{L}} \kappa_{-\mathbf{L}} \rangle}. \quad (4)$$

This final expression is our starting point to derive foreground quadratic estimators.

B. Foreground quadratic estimators

Following the argument above, anytime a field s (e.g. a foreground, the mask, the beam) has a non-zero bispectrum $\langle s_{\ell} s_{\mathbf{L}-\ell} s_{-\mathbf{L}} \rangle$, we can define the following linear

response¹:

$$f_{\ell, \mathbf{L}-\ell}^s = \frac{\langle s_{\ell} s_{\mathbf{L}-\ell} s_{-\mathbf{L}} \rangle}{\langle s_{\mathbf{L}} s_{-\mathbf{L}} \rangle} \quad (5)$$

such that

$$\langle s_{\ell} s_{\mathbf{L}-\ell} \rangle_{\text{at fixed } s_{\mathbf{L}}} = f_{\ell, \mathbf{L}-\ell}^s s_{\mathbf{L}} + \mathcal{O}((s_{\mathbf{L}})^2). \quad (6)$$

We can then define the ‘‘building block’’ foreground estimator, unbiased to first order in the foreground s , via $T_{\ell} T_{\mathbf{L}-\ell} / f_{\ell, \mathbf{L}-\ell}^s$. Finally, we inverse-variance weight these estimators for all ℓ , at fixed \mathbf{L} , to obtain the standard minimum variance quadratic estimator for s . In practice, evaluating this estimator requires knowledge of the foreground non-Gaussianity, through its bispectrum. We note however that no assumption about the foreground trispectrum is used.

C. Bias-hardened lensing quadratic estimators

In this section we review the derivation of the bias-hardened (BH) lensing estimator of [10, 16] for a single foreground, relate the noise of the BH estimator to the noise of the standard minimum variance QE, and generalize the procedure for an arbitrary number of foregrounds.

1. Hardening against a single foreground

In the presence of a single foreground s , the off-diagonal covariances of the observed temperature map take the form

$$\langle T_{\ell} T_{\mathbf{L}-\ell} \rangle = f_{\ell, \mathbf{L}-\ell}^{\kappa} \kappa_{\mathbf{L}} + f_{\ell, \mathbf{L}-\ell}^s s_{\mathbf{L}} \quad (7)$$

to lowest order in κ and s . In this equation the average is taken at fixed $\kappa_{\mathbf{L}}$ and $s_{\mathbf{L}}$, which will be implied from here on out for notational economy. Since both $\kappa_{\mathbf{L}}$ and $s_{\mathbf{L}}$ produce off-diagonal covariances, the standard quadratic estimators for κ and s acquire a bias. These biases can be written in the form [10]:

$$\begin{pmatrix} \langle \hat{\kappa}_{\mathbf{L}} \rangle \\ \langle \hat{s}_{\mathbf{L}} \rangle \end{pmatrix} = \begin{pmatrix} 1 & N_{\mathbf{L}}^{\kappa} \mathcal{R}_{\mathbf{L}} \\ N_{\mathbf{L}}^s \mathcal{R}_{\mathbf{L}} & 1 \end{pmatrix} \begin{pmatrix} \kappa_{\mathbf{L}} \\ s_{\mathbf{L}} \end{pmatrix} \quad (8)$$

where the response $\mathcal{R}_{\mathbf{L}}$ is defined to be

$$\mathcal{R}_{\mathbf{L}} = \int \frac{d^2\ell}{(2\pi)^2} \frac{f_{\ell, \mathbf{L}-\ell}^{\kappa} f_{\ell, \mathbf{L}-\ell}^s}{2C_{\ell}^{\text{tot}} C_{|\mathbf{L}-\ell|}^{\text{tot}}}. \quad (9)$$

We see that the bias to $\hat{\kappa}$ is proportional to s , and vice-versa. This allows us to create new bias-hardened estimators, which are the unique linear combinations of $\hat{\kappa}$ and

¹ We note that this expression only accounts for the non-Gaussianity of the sources. In principle, the sources are themselves lensed which results in an additional correlation and κ dependence of the mask. For simplicity, we ignore this effect in this work.

\hat{s} that null these biases

$$\begin{pmatrix} \hat{\kappa}_{\mathbf{L}}^{\text{BH}} \\ \hat{s}_{\mathbf{L}}^{\text{BH}} \end{pmatrix} = \begin{pmatrix} 1 & N_{\mathbf{L}}^{\kappa} \mathcal{R}_{\mathbf{L}} \\ N_{\mathbf{L}}^s \mathcal{R}_{\mathbf{L}} & 1 \end{pmatrix}^{-1} \begin{pmatrix} \hat{\kappa}_{\mathbf{L}} \\ \hat{s}_{\mathbf{L}} \end{pmatrix}. \quad (10)$$

By taking an average of Eq. (10), and inserting Eq. (8) on the right hand side, we see that bias-hardened estimators indeed satisfy $\langle \hat{\kappa}_{\mathbf{L}}^{\text{BH}} \rangle = \kappa_{\mathbf{L}}$ and $\langle \hat{s}_{\mathbf{L}}^{\text{BH}} \rangle = s_{\mathbf{L}}$.

2. Noise of the hardened lensing QE

By construction, the weights of the standard QE are chosen to minimize its variance. Since the BH estimator's weights differ from the standard QE's, bias-hardening comes at a price in noise, which we quantify in this section. We start with Eq. (10), the definition of the BH estimator, from which we compute the variance

$$\begin{aligned} \langle \hat{\kappa}_{\mathbf{L}}^{\text{BH}} \hat{\kappa}_{-\mathbf{L}}^{\text{BH}} \rangle &= \left[\langle \hat{\kappa}_{\mathbf{L}} \hat{\kappa}_{-\mathbf{L}} \rangle + (N_{\mathbf{L}}^{\kappa} \mathcal{R}_{\mathbf{L}})^2 \langle \hat{s}_{\mathbf{L}} \hat{s}_{-\mathbf{L}} \rangle \right. \\ &\quad \left. - 2N_{\mathbf{L}}^{\kappa} \mathcal{R}_{\mathbf{L}} \text{Re} \langle \hat{\kappa}_{\mathbf{L}} \hat{s}_{-\mathbf{L}} \rangle \right] / (1 - N_{\mathbf{L}}^{\kappa} N_{\mathbf{L}}^s \mathcal{R}_{\mathbf{L}}^2)^{-2}. \end{aligned} \quad (11)$$

In this equation we've used the identity $N_{\mathbf{L}}^{\kappa} = N_{-\mathbf{L}}^{\kappa}$, and we've assumed that $f_{\ell, \mathbf{L}-\ell}^s$ is even in both of its arguments, which forces both the noise of the source estimator and the response to be even functions of \mathbf{L} . This assumption is valid when modeling the field s as a collection of individual sources with profiles u_{ℓ} (since $u_{\ell} = u_{-\ell}$ for any physical emission profile), as in section III.

The cross-correlation of $\hat{\kappa}$ and \hat{s} is proportional to the response function, and is explicitly given by

$$\langle \hat{\kappa}_{\mathbf{L}} \hat{s}_{\mathbf{L}'} \rangle = (2\pi)^2 \delta_{\mathbf{L}+\mathbf{L}'}^D N_{\mathbf{L}}^{\kappa} N_{\mathbf{L}'}^s \mathcal{R}_{\mathbf{L}}. \quad (12)$$

Combining Eq. (11) and (12) results in the simple expression

$$N_{\mathbf{L}}^{\kappa, \text{BH}} = N_{\mathbf{L}}^{\kappa} (1 - N_{\mathbf{L}}^{\kappa} N_{\mathbf{L}}^s \mathcal{R}_{\mathbf{L}}^2)^{-1}. \quad (13)$$

We see that the noise cost for bias-hardening is completely determined by the determinant of the matrix in Eq. (10). The value of this determinant is plotted in Fig. 1 for the case of point source foregrounds ($f_{\ell, \mathbf{L}-\ell}^s = 1$). Nearly identical calculations yield the noise of the bias-hardened source estimator and the noise cross-correlation of the two bias-hardened estimators. The results are

$$\begin{aligned} N_{\mathbf{L}}^{s, \text{BH}} &= N_{\mathbf{L}}^s (1 - N_{\mathbf{L}}^{\kappa} N_{\mathbf{L}}^s \mathcal{R}_{\mathbf{L}}^2)^{-1} \\ \langle \hat{\kappa}_{\mathbf{L}}^{\text{BH}} \hat{s}_{\mathbf{L}}^{\text{BH}} \rangle &= -\langle \hat{\kappa}_{\mathbf{L}} \hat{s}_{\mathbf{L}} \rangle (1 - N_{\mathbf{L}}^{\kappa} N_{\mathbf{L}}^s \mathcal{R}_{\mathbf{L}}^2)^{-1}. \end{aligned} \quad (14)$$

3. Simultaneously hardening against n foregrounds

It is straightforward to generalize the bias-hardening procedure for n foreground fields $s_{1, \mathbf{L}}, \dots, s_{n, \mathbf{L}}$. In this case the observed map has off-diagonal covariances

$$\langle T_{\ell} T_{\mathbf{L}-\ell} \rangle = f_{\ell, \mathbf{L}-\ell}^{\kappa} \kappa_{\mathbf{L}} + f_{\ell, \mathbf{L}-\ell}^{s_1} s_{1, \mathbf{L}} + \dots + f_{\ell, \mathbf{L}-\ell}^{s_n} s_{n, \mathbf{L}} \quad (15)$$




Figure 1. Bias-hardening against point source foregrounds. The solid curve is the determinant of the matrix in Eq. (10), which is the ratio of the bias-hardened estimator's noise to the standard QE's noise. It is close to unity on most scales, implying that the noise penalty for bias hardening is minimal. The dashed curve is the ratio of the SNR of the point source estimator to the relative bias to $C_{\mathbf{L}}^{\kappa}$ from the source trispectrum. This ratio is roughly 10 at all scales. Thus in the regime where the bias to the standard QE is significant, the source estimator reconstructs the sources with a high SNR, and the source-hardened estimator can remove the bias due to sources with little noise cost. The spike is due to a zero crossing in the response $\mathcal{R}_{\mathbf{L}}$ at $L \sim 2500$.

to lowest order in κ and the foreground fields. Just as before, this average is taken at fixed $\kappa_{\mathbf{L}}, s_{1, \mathbf{L}}, \dots, s_{n, \mathbf{L}}$. First, we write down the standard minimum variance quadratic estimators for each of the fields:

$$\begin{aligned} \hat{X}_{\mathbf{L}} &= N_{\mathbf{L}}^X \int \frac{d^2 \ell}{(2\pi)^2} F_{\ell, \mathbf{L}-\ell}^X T_{\ell} T_{\mathbf{L}-\ell} \\ (N_{\mathbf{L}}^X)^{-1} &= \int \frac{d^2 \ell}{(2\pi)^2} F_{\ell, \mathbf{L}-\ell}^X f_{\ell, \mathbf{L}-\ell}^X \\ F_{\ell, \mathbf{L}-\ell}^X &= \frac{f_{\ell, \mathbf{L}-\ell}^X}{2C_{\ell}^{\text{tot}} C_{|\mathbf{L}-\ell|}^{\text{tot}}}, \end{aligned} \quad (16)$$

where $X = \kappa, s_1, \dots, s_n$. We then write down the $(n+1) \times (n+1)$ matrix that quantifies the biases to all the estimators, analogous to Eq. (8). Finally, we invert this matrix to obtain the bias-hardened estimators

$$\begin{pmatrix} \hat{\kappa}_{\mathbf{L}}^{\text{BH}} \\ \hat{s}_{1, \mathbf{L}}^{\text{BH}} \\ \vdots \\ \hat{s}_{n, \mathbf{L}}^{\text{BH}} \end{pmatrix} = \underbrace{\begin{pmatrix} 1 & N_{\mathbf{L}}^{\kappa} \mathcal{R}_{\mathbf{L}}^{s_1, \kappa} & \dots & N_{\mathbf{L}}^{\kappa} \mathcal{R}_{\mathbf{L}}^{s_n, \kappa} \\ N_{\mathbf{L}}^{s_1} \mathcal{R}_{\mathbf{L}}^{\kappa, s_1} & 1 & \dots & N_{\mathbf{L}}^{s_1} \mathcal{R}_{\mathbf{L}}^{s_1, s_n} \\ \vdots & \vdots & \ddots & \vdots \\ N_{\mathbf{L}}^{s_n} \mathcal{R}_{\mathbf{L}}^{s_n, \kappa} & N_{\mathbf{L}}^{s_n} \mathcal{R}_{\mathbf{L}}^{s_n, s_1} & \dots & 1 \end{pmatrix}^{-1}}_{\equiv M_{\mathbf{L}}^{-1}} \begin{pmatrix} \hat{\kappa}_{\mathbf{L}} \\ \hat{s}_{1, \mathbf{L}} \\ \vdots \\ \hat{s}_{n, \mathbf{L}} \end{pmatrix} \quad (17)$$

where the generalized response is defined to be

$$\mathcal{R}_{\mathbf{L}}^{X, Y} = \int \frac{d^2 \ell}{(2\pi)^2} \frac{f_{\ell, \mathbf{L}-\ell}^X f_{\ell, \mathbf{L}-\ell}^Y}{2C_{\ell}^{\text{tot}} C_{|\mathbf{L}-\ell|}^{\text{tot}}}. \quad (18)$$

Just as in the previous section, the response $\mathcal{R}_{\mathbf{L}}^{X,Y}$ is proportional to the cross-correlation of the two fields X and Y . When $X = Y$, the response is simply the inverse of the noise, which is manifested in Eq. (17) by the 1's on the diagonal. Writing the matrix $\mathbf{M}_{\mathbf{L}}$ as

$$\mathbf{M}_{\mathbf{L}} = \begin{pmatrix} 1 & \mathbf{w}_{\mathbf{L}}^T \\ \mathbf{v}_{\mathbf{L}} & \mathbf{A}_{\mathbf{L}} \end{pmatrix}, \quad (19)$$

where $\mathbf{A}_{\mathbf{L}}$ is a $n \times n$ matrix and $\mathbf{w}_{\mathbf{L}}, \mathbf{v}_{\mathbf{L}}$ are n -dimensional column vectors, we find that the noise of the BH lensing estimator is simply

$$N_{\mathbf{L}}^{\kappa \text{BH}} = \frac{\det(\mathbf{A}_{\mathbf{L}})}{\det(\mathbf{M}_{\mathbf{L}})} N_{\mathbf{L}}^{\kappa}. \quad (20)$$

A derivation of this equation is given in Appendix C.

III. SOURCE QE AND SOURCE-HARDENED LENSING QE

A. Halo model for sources

In the previous section we reviewed how to bias-harden against any number of foregrounds. The only input to this procedure is the linear response $f_{\ell, \mathbf{L}-\ell}^s$ of each foreground, which can be calculated if the foreground's bispectrum and power spectrum are known. To estimate these quantities we adopt a halo model, writing each foreground $s(\mathbf{x})$ as a sum of individual sources with profiles $u(\mathbf{x})$:

$$s(\mathbf{x}) = \sum_i s_i u(\mathbf{x} - \mathbf{x}_i), \quad (21)$$

where s_i and \mathbf{x}_i are the flux and position of the i 'th source, and the profile $u(\mathbf{x})$ is assumed to be source independent. These sources are assumed to be a Poisson sampling of the matter density field. Because they trace matter on large scales, the source field s is correlated with the true lensing convergence field. As we show below, this correlation causes a bias in the CMB lensing cross- and auto-correlations. With this model, the linear response takes the form

$$f_{\ell, \mathbf{L}-\ell}^s = \frac{\langle s_{\ell} s_{\mathbf{L}-\ell} s_{-\mathbf{L}} \rangle}{\langle s_{\mathbf{L}} s_{-\mathbf{L}} \rangle} = \frac{\langle s_i^3 \rangle}{\langle s_i^2 \rangle} \frac{u_{\ell} u_{\mathbf{L}-\ell}}{u_{\mathbf{L}}}, \quad (22)$$

where in the second equality we assumed that the sources are sparse, such that the bispectrum and power spectrum are both shot noise dominated, and the source clustering is negligible. In this case the linear response simplifies to a separable function of ℓ and $\mathbf{L} - \ell$, and is proportional to the ratio of the third to the second moment of the individual source fluxes.

We can choose the normalization of the source estimator \hat{s} to remove the constant factor $\langle s_i^3 \rangle / \langle s_i^2 \rangle$. Doing so makes \hat{s} an estimator of $s' \equiv s \langle s_i^2 \rangle / \langle s_i^3 \rangle$, which has linear response $f_{\ell, \mathbf{L}-\ell}^{s'} = u_{\ell} u_{\mathbf{L}-\ell} / u_{\mathbf{L}}$. This simplifies the

expression of the source estimator to

$$\hat{s}_{\mathbf{L}} = N_{\mathbf{L}}^s \int \frac{d^2 \ell}{(2\pi)^2} \frac{T_{\ell} T_{\mathbf{L}-\ell}}{2C_{\ell}^{\text{tot}} C_{|\mathbf{L}-\ell|}^{\text{tot}}} \frac{u_{\ell} u_{\mathbf{L}-\ell}}{u_{\mathbf{L}}} \quad (23)$$

$$(N_{\mathbf{L}}^s)^{-1} = \int \frac{d^2 \ell}{(2\pi)^2} \frac{1}{2C_{\ell}^{\text{tot}} C_{|\mathbf{L}-\ell|}^{\text{tot}}} \left[\frac{u_{\ell} u_{\mathbf{L}-\ell}}{u_{\mathbf{L}}} \right]^2,$$

which can be quickly evaluated using FFT. We emphasize that the weights and normalization of the BH estimator are left unchanged by this renormalization of \hat{s} . Therefore the bias-hardening procedure requires no knowledge of the source amplitude, and the normalization of the profile u is completely arbitrary. For this reason we will drop the primes from here-on-out, defining $f_{\ell, \mathbf{L}-\ell}^s = u_{\ell} u_{\mathbf{L}-\ell} / u_{\mathbf{L}}$.

B. Noise cost and importance of the non-Gaussianity

Consider a single source s . If s were Gaussian, it would be indistinguishable from map noise, and the source estimator \hat{s} would not be able to pick it out. Here, we show how the non-Gaussianity of the source map controls both the signal-to-noise (SNR) of the source estimator, and the bias from sources to CMB lensing. For simplicity let s be a Poisson distributed point source field ($f_{\ell, \mathbf{L}-\ell}^s = u_{\ell} = 1$). In this case the mean number of sources \bar{n} determines the non-Gaussianity; as \bar{n} goes to infinity, the statistics converges to that of a Gaussian.

The point source estimator is an unbiased (to lowest order) estimator of s , so the SNR per Fourier mode is simply :

$$\text{SNR}_{\mathbf{L}}^s = \langle \hat{s}_{\mathbf{L}} \hat{s}_{-\mathbf{L}} \rangle / N_{\mathbf{L}}^s \quad (24)$$

$$= \mathcal{T}^s / N_{\mathbf{L}}^s. \quad (25)$$

Thus the SNR for the source estimator scales with the scale-independent source trispectrum \mathcal{T}^s : it is larger for low source densities which corresponds to a higher non-Gaussianity. As the number of sources increases, the field becomes increasingly Gaussian, and can no longer be distinguished from Gaussian noise in the map. The expression above is valid for mildly non-Gaussian source field, as is the case in reality. For the very highly non-Gaussian case, which is not relevant for the CMB, the source non-Gaussianity contributes also to the noise, leading to a saturation of the SNR to a finite value as \bar{n} goes to zero.

Even if the source field and lensing convergence were independent, point sources would cause a bias to the standard QE due to its trispectrum:

$$\frac{\text{bias to } C_{\mathbf{L}}^{\kappa} \text{ from } \mathcal{T}^s}{N_{\mathbf{L}}^{\kappa}} = \frac{\mathcal{T}^s}{N_{\mathbf{L}}^{\kappa}} \left(N_{\mathbf{L}}^{\kappa} \int \frac{d^2 \ell}{(2\pi)^2} F_{\ell, \mathbf{L}-\ell}^{\kappa} \right)^2$$

$$= \mathcal{T}^s N_{\mathbf{L}}^{\kappa} \mathcal{R}_{\mathbf{L}}^2. \quad (26)$$

Here again, the source non-Gaussianity controls the bias to lensing via the $1/\bar{n}$ scaling. From equations (25) and (26), we see that the ratio of the SNR of the source estimator to the relative bias to lensing from the source

trispectrum is $(N_L^k N_L^s \mathcal{R}_L^2)^{-1}$. This ratio cancels out the non-Gaussianity of the sources and is around 10 for most scales of interest, as shown in Fig 1. This means that subtracting the lensing bias due to the source trispectrum will always cause only a small loss in lensing SNR. Indeed, as shown in Fig 1 (solid line), the lensing noise for the standard QE and the bias hardened estimators are very similar, meaning that subtracting the source bias causes a negligible increase in the lensing noise. This is consistent with the findings of [19].

Finally, the non-Gaussianity of the point sources also contributes an extra noise term in the lensing reconstruction, which we haven't included. However, this noise is only important when the bias to CMB lensing is also important. Thus, by making sure that the non-Gaussian point sources (and other foregrounds) do not significantly bias the lensing signal, we are making sure that their noise contribution is also small.

IV. BIASES TO CMB LENSING & FOREGROUND HARDENING

In this section we quantify the biases for three different bias-hardening schemes, and compare these results to the standard QE and Shear estimators. We consider two source components: point-sources $f_{\ell, L-\ell}^{\text{PS}} = 1$ and sources with tSZ-like profiles $f_{\ell, L-\ell}^{\text{tSZ}} = u_{\ell} u_{L-\ell} / u_L$. The profile u_{ℓ} is calculated by taking the square root of the tSZ power spectrum, which is measured from the simulations. Our three different bias-hardening schemes harden against one or both of these source components. The Point Source Hardened (PSH) estimator hardens against point sources; the Profile Hardened (PH) estimator hardens against a tSZ-like cluster profile, as described in App. D; and the Point source and Profile Hardened (PPH) estimator simultaneously hardens against both. We show the noise power spectra for these various estimators in Fig. 2, and performed various checks of our code in App. F.

A. Bias to CMB lensing from non-Gaussian foregrounds

The presence of a non-Gaussian foreground component s in the temperature map produces three bias terms to CMB lensing, referred to as primary, secondary and trispectrum terms [9, 10, 12, 13]. Indeed, if $T = T^{\text{CMB}} + s$, then the quadratic estimator \mathcal{Q} applied to T can be expanded bilinearly as:

$$\mathcal{Q}[T, T] = \mathcal{Q}[T^{\text{CMB}}, T^{\text{CMB}}] + 2\mathcal{Q}[T^{\text{CMB}}, s] + \mathcal{Q}[s, s], \quad (27)$$

assuming the quadratic estimator is symmetric (or has been symmetrized in its arguments).

In cross-correlation with a mass tracer g , this simply leads to the bispectrum bias:

$$\langle g\mathcal{Q} \rangle = \underbrace{\langle g\kappa_{\text{CMB}} \rangle}_{\text{lensing signal}} + \underbrace{\langle g\mathcal{Q}[s, s] \rangle}_{\text{bispectrum bias}}. \quad (28)$$




Figure 2. Noises of the Point Source Hardened (PSH), Profile Hardened (PH), Point source and Profile Hardened (PPH), Shear, and standard quadratic (QE) estimators for $\ell_{\text{max}, T} = 3500$. In black is the lensing signal C_L^k . Hardening against a single source (PSH, PH) marginally increases the noise over QE, whereas hardening against multiple sources (PPH) results in a larger noise penalty. We also note that the noise cost of bias-hardening increases with the size of the source profile u_{ℓ} , which is discussed further in Appendix E.

For the auto-correlation, more terms arise:

$$\begin{aligned} \langle \mathcal{Q}\mathcal{Q} \rangle &= \underbrace{\langle \kappa_{\text{CMB}} \kappa_{\text{CMB}} \rangle}_{\text{lensing signal}} \\ &+ 2 \underbrace{\langle \mathcal{Q}[T^{\text{CMB}}, T^{\text{CMB}}] \mathcal{Q}[s, s] \rangle}_{\text{Primary bispectrum bias}} \\ &+ 4 \underbrace{\langle \mathcal{Q}[T^{\text{CMB}}, s] \mathcal{Q}[T^{\text{CMB}}, s] \rangle}_{\text{Secondary bispectrum bias}} \\ &+ \underbrace{\langle \mathcal{Q}[s, s] \mathcal{Q}[s, s] \rangle}_{\text{Trispectrum bias}}. \end{aligned} \quad (29)$$

The primary and secondary bias terms are due to the non-Gaussianity of the foreground and its correlation with the true CMB lensing signal. They are integrals of the $\kappa_{\text{CMB}} s s$ bispectrum. In particular, predicting them requires knowing the redshift distribution of the foreground. On the other hand, the trispectrum bias is present whether or not the sources are correlated with CMB lensing, and only depend on the statistics (the trispectrum) of the projected 2d foreground map.

In what follows, we compute these biases for the various quadratic lensing estimators of interest. We follow the approach of [13], using the simulations from [20].

B. Effectiveness of foreground hardening: expectations

Here we explain which of these bias terms are nulled with bias-hardening. The only knowledge of foregrounds used in the bias hardening process is the shape of the foreground bispectrum and power spectrum, through f^s . In particular, the procedure does not assume anything about the $ss\kappa_{\text{CMB}}$ bispectrum, which causes the primary and secondary biases, nor about the $ssss$ trispectrum, which causes the trispectrum bias. Why would bias hardening help reducing foreground biases at all then? The answer is somewhat subtle.

The bias hardening procedure nulls the linear response of the lensing estimator to the foreground s , i.e. the bispectrum $\langle \mathcal{Q}^{\text{BH}}[s, s] s \rangle$. This is very similar but not identical to the primary bispectrum bias $\langle \mathcal{Q}^{\text{BH}}[s, s] \kappa_{\text{CMB}} \rangle$. The two bispectra are all the more similar as the foreground map and the true lensing map are highly correlated. Thus in principle, the primary bispectrum bias is only partially reduced by the bias-hardening. In practice though, as we show below with simulated foregrounds, this reduction is very large.

As for the secondary bispectrum bias, there is no reason *a priori* why the bias hardening procedure would reduce it. Indeed, we show below that the various lensing estimators have a similar level of secondary bias.

The case of the trispectrum bias is interesting. In general, if $\mathcal{T}_{\ell_1, \ell_2, \ell_3, \ell_4}$ is the source trispectrum, the trispectrum bias to lensing is given by

$$\langle \mathcal{Q}[s, s] \mathcal{Q}[s, s] \rangle_c = N_{\mathbf{L}}^2 \int \frac{d^2 \ell}{(2\pi)^2} F_{\ell, \mathbf{L}-\ell} \int \frac{d^2 \ell'}{(2\pi)^2} F_{\ell', -\mathbf{L}-\ell'} \mathcal{T}_{\ell, \mathbf{L}-\ell, \ell', -\mathbf{L}-\ell'}.$$
 (30)

If the lensing weights of the bias-hardened estimator are denoted as $F_{\ell, \mathbf{L}-\ell}^{\text{BH}}$, then the zero linear response to the foreground can be translated as:

$$\int \frac{d^2 \ell}{(2\pi)^2} F_{\ell, \mathbf{L}-\ell}^{\text{BH}} f_{\ell, \mathbf{L}-\ell}^s = 0.$$
 (31)

This does not null the trispectrum bias in general. However, it does in the case of Poisson sources with identical profiles. Indeed, in this case, the trispectrum $\mathcal{T}_{\ell_1, \ell_2, \ell_3, \ell_4} = \bar{n} \langle s^4 \rangle u_{\ell_1} u_{\ell_2} u_{\ell_3} u_{\ell_4}$ is separable, such that the trispectrum bias becomes:

$$\bar{n} \langle s^4 \rangle N_{\mathbf{L}}^2 \left[\int \frac{d^2 \ell}{(2\pi)^2} F_{\ell, \mathbf{L}-\ell}^{\text{BH}} u_{\ell} u_{\mathbf{L}-\ell} \right]^2.$$
 (32)

Since the response $f_{\ell, \mathbf{L}-\ell}^s = u_{\ell} u_{\mathbf{L}-\ell} / u_{\mathbf{L}} \propto u_{\ell} u_{\mathbf{L}-\ell}$, we see that Eq. (31) implies that the trispectrum bias is exactly nulled. However, it is no longer generally true if the sources are clustered, or if there is a distribution of profile sizes or shapes. Because our profile hardening assumes a single fiducial source profile, the trispectrum bias is also not exactly nulled when the foreground sources have a range of profiles, as is the case for the tSZ.

In summary, the PH estimator should significantly reduce the tSZ trispectrum bias, while the PSH estimator

should exactly null the point source trispectrum. Both should largely reduce the primary bias, with no *a priori* expectation for the secondary bias.

For bias hardening against more complex foregrounds, such as clustered sources, the effectiveness of this simple bias hardening is not guaranteed. However, if the statistical properties of the source distribution is known, it is nonetheless possible to extend our formalism to harden against non-linear clustering. This has been explored in the context of line intensity mapping in [21], and we believe that a similar treatment could be helpful for mitigating biases to CMB lensing.

C. Effectiveness of foreground hardening: simulations

To quantify the biases of each estimator for each of the extragalactic foregrounds, we use simulated maps [20] of the CIB, tSZ, kSZ, radio point sources, and κ_{CMB} at 148 GHz. We rescale each of the foregrounds (0.38 for CIB, 0.7 for tSZ, 0.82 for kSZ, 1.1 for radio PS) to match [22], then subtract the mean from each foreground, and finally mask each foreground using a matched filter for point sources with a flux cut of 5 mJy and mask patch radius $3'$. The resulting power spectra are plotted in Appendix A.

For CMB measurements, we consider an experiment similar to Simons Observatory [5], with a white noise level of $6 \mu\text{K-arcmin}$ and a Gaussian beam with full-width at half maximum of $1.4'$.

As in [13], we re-weight the halos from [20] to match the redshift distribution of the LSST gold sample ($dn/dz \propto (z/0.24)^2 e^{-z/0.24} / 0.48$) [23] and obtain a projected galaxy number density count δ_g . Altogether, a set of these maps (foregrounds, κ_{CMB} , and δ_g) allow us to calculate the biases to the cross-correlation of CMB lensing with galaxy number density counts $C_L^{\kappa \delta_g}$, as well as the primary, secondary, and trispectrum biases to the lensing auto-correlation $C_L^{\kappa \kappa}$. We run the estimators on 81 flat square cutouts from these maps, obtaining the biases plotted in Fig. 3 and 4.

For the cross-correlation of CMB lensing with galaxy counts (Fig. 3), we find that the standard minimum variance quadratic estimator acquires a $\sim 10\%$ bias from the CIB and tSZ for reasonable values of $\ell_{\text{max}, \text{T}}$. Hardening against a single source component with a tSZ-like profile (PH) effectively nulls the bias from both tSZ and kSZ. Hardening against just point sources (PSH) is complementary, effectively nulling the radio point source and CIB biases. When simultaneously hardening against both tSZ-like profiles and point sources (PPH), we're able to null the biases from all extragalactic foregrounds, creating a highly robust estimator of the cross-correlation.

For the CMB lensing auto-correlation we again find that the standard QE acquires a $\sim 10\%$ bias, primarily due to the CIB and tSZ. This is consistent with the results found in [10]. In Fig. 4 we see that the PH estimator nulls the primary bias from tSZ, and significantly reduces the remaining primary and trispectrum biases. Likewise the PSH estimator nulls the primary bias from point sources,




Figure 3. Relative bias on the cross with LSST at $\ell_{\max, T} = 3500$. The gray boxes are the statistical uncertainties when reconstructing $C_L^{\kappa_{\delta g}}$ with the standard QE. We find that hardening against point sources (PSH) nulls or dramatically reduces the bias from radio point sources and CIB, whereas hardening against a tSZ-like profile (PH) nulls or significantly reduces the bias from tSZ and kSZ. Hardening against both (PPH) effectively nulls the bias from all foregrounds, resulting in a bias that is roughly an order of magnitude smaller than the bias to the standard QE.

while also significantly reducing remaining primary and trispectrum biases. The PPH estimator essentially nulls all primary biases, and reduces the trispectrum bias to a level comparable to Shear. As discussed in the previous section, there is no reason to expect a significant reduction to the secondary biases from bias-hardening. In practice we find that the secondary bias to bias-hardened estimators is slightly worse than (but comparable to) the bias to the standard QE. We note that the secondary bias is negative, whereas the bias from the trispectrum is positive. This results in a cancellation in the overall bias to the lensing power, allowing higher values of $\ell_{\max, T}$ than one might expect from looking at the secondary and trispectrum biases alone.

As a figure of merit for the performance of each estimator, we use the signal-to-noise ratio (SNR) of the lensing amplitude A_{lens} , defined for each L as $C_L^{\kappa, m} / C_L^{\kappa}$, where $C_L^{\kappa, m}$ and C_L^{κ} are the measured and true lensing power spectra respectively. By inverse variance weighting the measurement of A_{lens} for each L , we obtain an estimator for the lensing amplitude:

$$\hat{A}_{\text{lens}} = \int \frac{d^2 \mathbf{L}}{(2\pi)^2} \frac{C_L^{\kappa, m}}{C_L^{\kappa}} \frac{(C_L^{\kappa})^2}{\sigma_L^2} \bigg/ \int \frac{d^2 \mathbf{L}}{(2\pi)^2} \frac{(C_L^{\kappa})^2}{\sigma_L^2} \quad (33)$$

with $\sigma_L^2 = 2(C_L^{\kappa} + N_L)^2$.

Since A_{lens} has a fiducial value of 1, the SNR and bias of

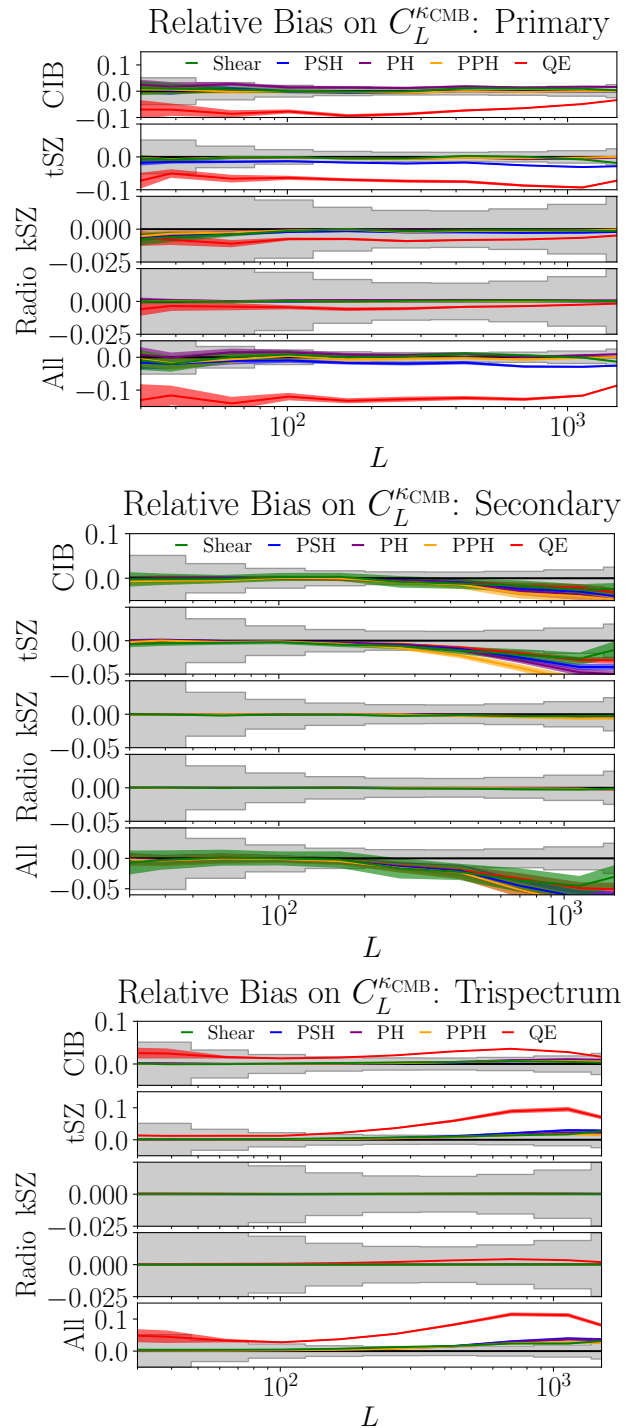


Figure 4. Bias to the lensing auto-correlation at $\ell_{\max, T} = 3500$. The gray boxes are the statistical uncertainties when reconstructing C_L^{κ} with the standard QE. Hardening against point sources (PSH) significantly reduces the primary and trispectrum bias from radio point sources and CIB, whereas hardening against a tSZ-like profile significantly reduces these biases from tSZ. Simultaneously hardening against both (PPH) effectively nulls the primary bias. Bias-hardening has no significant effect on the secondary bias.

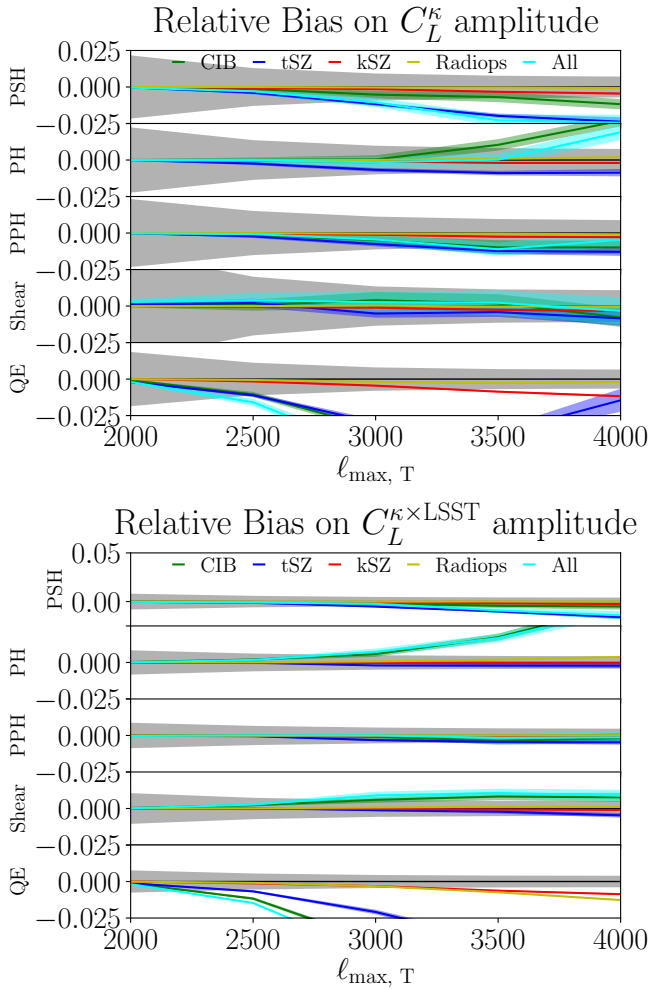


Figure 5. Relative bias on the lensing amplitude and the cross-correlation with LSST as a function of $\ell_{\max, T}$. In gray is the 1σ error for each estimator $1/\text{SNR}(A_{\text{lens}})$. CIB and tSZ are the main sources of bias to the standard QE. The Point Source Hardened (PSH) estimator drastically reduces the bias from CIB, whereas the Profile Hardened (PH) estimator reduces the bias from tSZ. The Point source and Profile Hardened (PPH) estimator reduces the bias from both tSZ and CIB to a level comparable to Shear while at a lower noise cost.

A_{lens} are

$$\text{SNR}^2(A_{\text{lens}}) = \int \frac{d^2\mathbf{L}}{(2\pi)^2} \frac{(C_L^\kappa)^2}{\sigma_L^2}$$

$$\text{bias}(A_{\text{lens}}) = \int \frac{d^2\mathbf{L}}{(2\pi)^2} \frac{\text{bias}(C_L^\kappa)}{C_L^\kappa} \frac{(C_L^\kappa)^2}{\sigma_L^2} \bigg/ \int \frac{d^2\mathbf{L}}{(2\pi)^2} \frac{(C_L^\kappa)^2}{\sigma_L^2}. \quad (34)$$

Similar definitions for the amplitude of the CMB-LSST cross-correlation can be made by replacing the lensing power C_L^κ with $C_L^{\kappa\delta_g}$ and by setting $\sigma_L^2 = (C_L^\kappa + N_L^\kappa)(C_L^{\delta_g} + N_L^{\delta_g}) + (C_L^{\kappa\delta_g})^2$ in the equations above. We show the best SNR achievable in auto and cross-correlation for the various estimators, while maintaining the foreground bias below 1σ , in Tab. I.

For the auto-correlation, we find that the Profile Hardened (PH) estimator achieves the highest SNR be-




Figure 6. Total signal-to-noise of the lensing amplitude (top) and the cross correlation with LSST (bottom). The black markers denote the highest $\ell_{\max, T}$ where the bias is less than 1σ . We find that the Profile Hardened (PH) estimator reconstructs the lensing amplitude with the highest signal to noise: a $\sim 50\%$ improvement over the standard QE. We find that the Point Source Hardened (PSH) and Point source and Profile Hardened (PPH) reconstruct the cross with similar SNRs: a $\sim 50\%$ improvement over the standard QE. The Hybrid estimator is a combination of Shear and QE, and is described in greater detail in Appendix B. We note that the color-scale saturates at 2σ , and that the bias to the standard QE can be $> 10\sigma$ at $\ell_{\max, T} \gtrsim 3000$.

fore $\text{bias}(A_{\text{lens}})$ becomes larger than the noise $\sigma = 1/\text{SNR}(A_{\text{lens}})$, as shown in Fig. 6. This is a $\sim 50\%$ improvement over the standard QE. The PH estimator's 1σ bias threshold occurs at $\ell_{\max, T} \sim 3700$. From Fig. 5, we see that this high value of $\ell_{\max, T}$ is due to a cancellation in the tSZ and CIB biases. A more conservative stopping point would be when any one of the foregrounds crosses the 1σ threshold, which would be around $\ell_{\max, T} \sim 3200$ for the PH estimator. We note that even at this lower $\ell_{\max, T}$, the PH estimator still has the highest SNR and is still a $\sim 30\%$ improvement over QE, which crosses the 1σ bias threshold at $\ell_{\max, T} \sim 2400$.

For the cross-correlation, the PSH and PPH achieve similar SNRs at the 1σ bias threshold, improving upon

QE by $\sim 40\%$. We note that the PPH has a stable bias across the entire range of realistic $\ell_{\max, T}$'s, making it a conservative yet competitive option for reconstructing the cross-correlation.

We conclude by noting that for an experiment with $\ell_{\max, T} \sim 3000$, the PH and PPH estimators reconstruct the auto- and cross-correlations respectively with the highest SNR while achieving a sub-percent level bias, as shown in Tab. II.

	SNR in auto ($\ell_{\max, T}$)	SNR in cross ($\ell_{\max, T}$)
PSH	104 (2936)	217 (3175)
PH	126 (3705)	189 (2822)
PPH	101 (3408)	215 (4000)
Shear	92 (4000)	158 (2792)
QE	81 (2375)	150 (2188)

Table I. Comparison of the total signal-to-noise ratio on the amplitude of the CMB lensing auto-spectrum (first column) and cross-spectrum with LSST-like galaxies (second column). In all cases, the maximum temperature multipole $\ell_{\max, T}$ is selected to guarantee a foreground bias smaller than 1σ .

	Rel. bias to auto (cross) [%]	Rel. bias to auto (cross) [σ]	SNR on C_L^κ ($C_L^{\kappa \times \text{LSST}}$)
PSH	-1.1 (-0.2)	-1.14 (-0.46)	108 (209)
PH	-0.1 (0.7)	-0.15 (1.47)	102 (202)
PPH	-0.4 (0.1)	-0.39 (0.01)	89 (184)
Shear	0.3 (0.9)	0.19 (1.55)	74 (173)
QE	-4.9 (-5.0)	-6.15 (-11.16)	125 (223)

Table II. Relative bias (in both % and σ units) and SNR for each estimator at $\ell_{\max, T} = 3000$. The PH and PPH estimators reconstruct the auto- and cross-correlations respectively with the highest SNR while achieving a sub-percent level bias.

V. DISCUSSION AND CONCLUSIONS

In this study we revisited foreground-hardened lensing quadratic estimators, described how to harden against arbitrary profiles, and extended the formalism to simultaneously harden against an arbitrary number of extragalactic foregrounds. We showed that the point source hardened estimator (PSH) reduces not only the bias from radio point sources, but also from the cosmic infrared background, the thermal and the kinematic Sunyaev-Zel'dovich effects.

Motivated by the extended nature of tSZ clusters, we modify the point source hardening technique to deproject the lensing bias from uncorrelated clusters (PH estimator), or simultaneously point sources and extended clusters (PPH estimator). We propose a simple and sufficient method to approximate the input cluster profile from the data itself, by considering the observed tSZ power spectrum from the same data.

The signal-to-noise cost from deprojecting the point source bias is small ($\sim 15\%$ for Point Source Harden-

ing), and is more than compensated by enabling the use of higher multipoles: overall, the SNR on the lensing power spectrum is increased by $\sim 56\%$ when using Profile Hardening (PH), while for cross-correlations, both PSH and PPH provide a $\sim 45\%$ increase in SNR compared to QE. Part of the improvement of PH and PPH estimators comes from the increase in the maximum multipole that can be included in the lensing reconstruction.

The PSH, PH and PPH estimators improve the lensing SNR over the shear and hybrid shear estimators of [13], although at the cost of a slightly larger bias in cross-correlation. They also provide consistency checks for the standard quadratic estimator, having a smaller foreground bias for a given ℓ_{\max} in temperature. Overall, these estimators outperform the standard lensing quadratic estimator both in auto and cross-correlation, both in terms of the lensing SNR and in terms of the foreground bias. We recommend their use in the CMB lensing analyses of upcoming temperature-dominated CMB data such as the Simons Observatory. Future surveys that are dominated by polarization reconstruction such as CMB-S4 [6], will still benefit from improved robustness, since temperature reconstruction will contribute a non-trivial fraction to the SNR. We note that a realistic analysis will require modeling of the higher order biases $N^{(i)}$ [19] that appear in lensing reconstruction and will likely be non-negligible for the next generation of surveys.

Finally, foreground hardening can and should be implemented in conjunction with other methods of foreground mitigation, such as masking [9] and multi-frequency cleaning in one or both legs of the estimator. We shall report on the optimal combination of these methods in an upcoming paper.

ACKNOWLEDGMENTS

We thank Anthony Challinor, Colin Hill, Antony Lewis, Mathew Madhavacheril, Blake Sherwin, Omar Darwish, Alex van Engelen, Neelima Sehgal, David Spergel, Martin White and the participants of the Sussex CMB lensing workshop for all their feedback. We give a special thanks to Zoom for facilitating this collaboration throughout the COVID-19 pandemic. 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.

Appendix A: Foreground power spectra

Our processing of the Sehgal simulations follows [13]. After masking the sources detected in the foreground maps at 5 mJy, which are detected at about 5σ given the noise level assumed, the power spectrum of each foreground at 148 GHz is shown in Fig. 7.




Figure 7. Foreground power spectra at 148 GHz after being masked with a match-filter with a flux cut of 5 mJy and a mask patch radius of $3'$. We assume a $6 \mu\text{K-arcmin}$ noise, and a $1.4'$ beam.

Appendix B: Hybrid estimator

The hybrid estimator is defined to be

$$\hat{\kappa}_L^H = N_L^H \left(\hat{\kappa}_L^{\text{QE}} / N_L^{\text{QE}} + \hat{\kappa}_L^{\text{S}} / N_L^{\text{S}} \right) \quad (\text{B1})$$

where the standard estimator $\hat{\kappa}_L^{\text{QE}}$ is calculated using $\ell_{\min, \text{T}} \leq \ell \leq 2000$ and the shear estimator $\hat{\kappa}_L^{\text{S}}$ is calculated using $2000 \leq \ell \leq \ell_{\max, \text{T}}$. The noise of the hybrid estimator is simply

$$N_L^H = \frac{1}{1/N_L^{\text{QE}} + 1/N_L^{\text{S}}}.$$

assuming that shear and QE are uncorrelated. The bias B_L to the lensing signal is defined by

$$C_L^{\text{QE}} = C_L^\kappa + N_L^{\text{QE}} + B_L^{\text{QE}} \quad (\text{B2})$$

$$C_L^{\text{S}} = C_L^\kappa + N_L^{\text{S}} + B_L^{\text{S}}. \quad (\text{B3})$$

Assuming that the noises of $\hat{\kappa}_L^{\text{QE}}$ and $\hat{\kappa}_L^{\text{S}}$ are uncorrelated, the bias to the Hybrid estimator is

$$B_L^H = (N_L^H)^2 \left(\frac{B_L^{\text{QE}}}{(N^{\text{QE}})^2} + \frac{B_L^{\text{S}}}{(N^{\text{S}})^2} \right). \quad (\text{B4})$$

Here B_L^{QE} is the bias when the standard quadratic estimator is calculated using $\ell_{\min, \text{T}} \leq \ell \leq 2000$. We approximate B_L^{S} as the bias from Shear when calculated with $\ell_{\min, \text{T}} \leq \ell \leq \ell_{\max, \text{T}}$. This approximation should be OK since the contribution from $\ell_{\min, \text{T}} \leq \ell \leq 2000$ is negligible. The bias to the hybrid estimator's cross-correlation with LSST is calculated by inverse noise weighting the biases to QE and Shear.

Appendix C: Derivation of Eq. (20)

We start with

$$M_L = \begin{pmatrix} 1 & N_L^\kappa \mathcal{R}_L^{\kappa, s_1} & \dots & N_L^\kappa \mathcal{R}_L^{\kappa, s_n} \\ N_L^{s_1} \mathcal{R}_L^{s_1, \kappa} & 1 & \dots & N_L^{s_1} \mathcal{R}_L^{s_1, s_n} \\ \vdots & \vdots & \ddots & \vdots \\ N_L^{s_n} \mathcal{R}_L^{s_n, \kappa} & N_L^{s_n} \mathcal{R}_L^{s_n, s_1} & \dots & 1 \end{pmatrix} \equiv \begin{pmatrix} 1 & \mathbf{w}_L^T \\ \mathbf{v}_L & \mathbf{A}_L \end{pmatrix}, \quad (\text{C1})$$

where \mathbf{A}_L is a $n \times n$ matrix and $\mathbf{w}_L, \mathbf{v}_L$ are n -dimensional column vectors. By splitting the matrix into blocks, we can express its inverse and determinant as

$$M_L^{-1} = \begin{pmatrix} (1 - \mathbf{w}_L^T \mathbf{A}_L^{-1} \mathbf{v}_L)^{-1} & -(1 - \mathbf{w}_L^T \mathbf{A}_L^{-1} \mathbf{v}_L)^{-1} \mathbf{w}_L^T \mathbf{A}_L^{-1} \\ -\mathbf{A}_L^{-1} \mathbf{v}_L (1 - \mathbf{w}_L^T \mathbf{A}_L^{-1} \mathbf{v}_L)^{-1} & \mathbf{A}_L^{-1} + \mathbf{A}_L^{-1} \mathbf{v}_L (1 - \mathbf{w}_L^T \mathbf{A}_L^{-1} \mathbf{v}_L)^{-1} \mathbf{w}_L^T \mathbf{A}_L^{-1} \end{pmatrix} \quad (\text{C2})$$

$$\det(M_L) = \det(\mathbf{A}_L) (1 - \mathbf{w}_L^T \mathbf{A}_L^{-1} \mathbf{v}_L). \quad (\text{C3})$$

Using this notation, the bias-hardened lensing estimator takes the form

$$\hat{\kappa}_L^{\text{BH}} = \frac{\hat{\kappa}_L - \mathbf{w}_L^T \mathbf{A}_L^{-1} \hat{\mathbf{s}}_L}{1 - \mathbf{w}_L^T \mathbf{A}_L^{-1} \mathbf{v}_L}, \quad (\text{C4})$$

where $\hat{\mathbf{s}}_L \equiv (\hat{s}_{1,L}, \hat{s}_{2,L}, \dots, \hat{s}_{n,L})^T$. The variance of the bias-hardened lensing estimator is simply

$$\begin{aligned} \langle \hat{\kappa}_L^{\text{BH}} \hat{\kappa}_{-L}^{\text{BH}} \rangle &= (1 - \mathbf{w}_L^T \mathbf{A}_L^{-1} \mathbf{v}_L)^{-2} \left[\langle \hat{\kappa}_L \hat{\kappa}_{-L} \rangle + \sum_{ij} \sum_{ab} (\mathbf{w}_L)_i (\mathbf{A}_L^{-1})_{ij} (\mathbf{w}_L)_a (\mathbf{A}_L^{-1})_{ab} \langle \hat{s}_{j,-L} \hat{s}_{b,L} \rangle \right. \\ &\quad \left. - 2 \sum_{ij} (\mathbf{w}_L)_i (\mathbf{A}_L^{-1})_{ij} \text{Re} \langle \hat{\kappa}_{-L} \hat{s}_{j,L} \rangle \right], \end{aligned} \quad (\text{C5})$$

from which we get the noise

$$\begin{aligned} N_L^{\kappa \text{BH}} &= (1 - \mathbf{w}_L^T \mathbf{A}_L^{-1} \mathbf{v}_L)^{-2} \left[N_L^\kappa + \sum_{ij} (\mathbf{w}_L)_i (\mathbf{A}_L^{-1})_{ij} N_L^{s_j} \sum_{ab} (\mathbf{w}_L)_a (\mathbf{A}_L^{-1})_{ab} N_L^{s_b} \mathcal{R}_L^{s_b, s_j} \right. \\ &\quad \left. - 2 N_L^\kappa \sum_{ij} (\mathbf{w}_L)_i (\mathbf{A}_L^{-1})_{ij} N_L^{s_j} \mathcal{R}_L^{\kappa, s_j} \right]. \end{aligned} \quad (\text{C6})$$

In the past two equations we've assumed that the noises and responses are even functions of \mathbf{L} , which is true so long that the sources are modeled as halos with profiles u_L . Let's focus on the second term of the RHS of the equation above. Note that $N_L^{s_b} \mathcal{R}_L^{s_b, s_j}$ is the b 'th component of the j 'th column vector of \mathbf{A}_L . Therefore

$$\sum_b (\mathbf{A}_L^{-1})_{ab} N_L^{s_b} \mathcal{R}_L^{s_b, s_j} = \delta_{aj}^K. \quad (\text{C7})$$

From this we find

$$\sum_{ab} (\mathbf{w}_L)_a (\mathbf{A}_L^{-1})_{ab} N_L^{s_b} \mathcal{R}_L^{s_b, s_j} = \sum_a (\mathbf{w}_L)_a \sum_b (\mathbf{A}_L^{-1})_{ab} N_L^{s_b} \mathcal{R}_L^{s_b, s_j} = \sum_a (\mathbf{w}_L)_a \delta_{aj}^K = (\mathbf{w}_L)_j. \quad (\text{C8})$$

Recall that the j 'th component of \mathbf{w}_L is just $N_L^\kappa \mathcal{R}_L^{\kappa, s_j}$. Plugging this result back into our expression for the noise gives

$$\begin{aligned} N_L^{\kappa \text{BH}} &= (1 - \mathbf{w}_L^T \mathbf{A}_L^{-1} \mathbf{v}_L)^{-2} \left[1 - \sum_{ij} (\mathbf{w}_L)_i (\mathbf{A}_L^{-1})_{ij} \underbrace{N_L^{s_j} \mathcal{R}_L^{\kappa, s_j}}_{=(\mathbf{w}_L)_j} \right] N_L^\kappa \\ &= (1 - \mathbf{w}_L^T \mathbf{A}_L^{-1} \mathbf{v}_L)^{-1} N_L^\kappa \\ &= \frac{\det(\mathbf{A}_L)}{\det(M_L)} N_L^\kappa. \end{aligned} \quad (\text{C9})$$

Appendix D: tSZ-like profile

If the tSZ foreground map is modeled as $s_\ell = \sum_i s_i u_\ell$, then the tSZ power spectrum is proportional to $|u_\ell|^2$. Recall that the bias-hardened estimator is insensitive to the normalization of the profile. Therefore we take the square root of the tSZ power spectrum (measured from the simulations) as our tSZ-like profile, which is plotted in Fig. 8. In practice, determining the exact profile from the measured tSZ power spectrum may be challenging. However, we have checked that the residual biases from tSZ and other foregrounds vary slowly when the assumed profile is changed, such that a precise knowledge of the profile is not needed. This is encouraging, and suggests that this method will be useful in practice.




Figure 8. tSZ-like profile u_ℓ . For $\ell \lesssim 60$ we approximate the profile as a flat curve.

Appendix E: Noise price of bias-hardening increases with profile size

We found that the noise price for hardening against a tSZ-like profile is larger than hardening against point sources. To gain an intuition for why this is the case, we plot the noise when hardening against a single Gaussian profile $u_\ell = e^{-\sigma^2 \ell^2 / 2}$ for different values of σ in Fig. 9. We find that the noise increases with σ ; that is, the larger the profile the larger the cost in noise. As shown in Fig. 1, in the regime where the response is small the noise cost goes as \mathcal{R}_L^2 . The response $\mathcal{R}_L \propto \int d^2 \ell f_{\ell, L-\ell}^s f_{\ell, L-\ell}^k / C_\ell^{\text{tot}} C_{|L-\ell|}^{\text{tot}}$ is larger when the linear response to sources $f_{\ell, L-\ell}^s = e^{\sigma^2 L^2 / 2} e^{-\sigma^2 \ell^2 / 2} e^{-\sigma^2 |L-\ell|^2 / 2}$ looks more like $f_{\ell, L-\ell}^k$. For point sources $f_{\ell, L-\ell}^s$ is flat, however, as the size of the profile gets larger $f_{\ell, L-\ell}^s$ looks slightly more similar to the linear response to lensing, making it more difficult for the source-hardened estimator to remove the sources, and results in a higher noise price.




Figure 9. **Shaded (blue to yellow) region:** Noise when bias hardening against a single Gaussian profile $f_{\ell, L-\ell}^s = e^{\sigma^2 L^2 / 2} e^{-\sigma^2 \ell^2 / 2} e^{-\sigma^2 |L-\ell|^2 / 2}$ for $\sigma \in [0', 2.4']$. σ is monotonically increasing from blue to yellow. Note that $\sigma = 0'$ corresponds to hardening against point sources (PSH). **Red:** Noise of the standard QE.

Appendix F: Noise and cross-correlation checks

As a simple pipeline check we verify our calculation of the noise and response when bias-hardening against point sources (PSH). These checks are shown in the plots below.




Figure 10. Verification of the noises of each estimator. The auto-correlations (blue) of the estimators $\hat{\kappa}$, \hat{s} , and $\hat{\kappa}^{\text{PSH}}$ when run on a Gaussian random field with power spectrum C_L^{tot} . In black are the theory curves.




Figure 11. Cross-correlating the true convergence κ with $\hat{\kappa}$, \hat{s} , and $\hat{\kappa}^{\text{PSH}}$ when run on a lensed CMB map. We see that both the standard QE $\hat{\kappa}$ and the point source hardened estimator $\hat{\kappa}^{\text{PSH}}$ recover the lensing signal C_L^κ when correlated with the true convergence. The cross correlation of the point source estimator \hat{s} with the true lensing convergence gives $N_L^s C_L^\kappa \mathcal{R}_L$, which is verified in the middle plot. The blue points are positive, whereas the red points are negative. The response has a zero crossing at $L \sim 2500$.




Figure 12. Cross-correlating the true source map s with $\hat{\kappa}$, \hat{s} , and $\hat{\kappa}^{\text{PSH}}$ when run on the true point source map s . The cross-correlation of $\hat{\kappa}$ with s gives $N_L^s C_L^s \mathcal{R}_L$, which is verified in the plot on the left. \hat{s} recovers the source signal C_L^s when cross correlated with the true source map. Since the point-source hardened estimator is designed to have zero response to point sources, the cross-correlation of $\hat{\kappa}^{\text{PSH}}$ and s should be zero, which is consistent with the result found in the rightmost plot.

-
- [1] A. Lewis and A. Challinor, *Phys. Rep.* **429**, 1 (2006).
 - [2] D. Hanson, A. Challinor, and A. Lewis, *General Relativity and Gravitation* **42**, 2197 (2010).
 - [3] S. W. Henderson *et al.*, *Journal of Low Temperature Physics* **184**, 772 (2016).
 - [4] B. A. Benson *et al.*, **9153**, 91531P (2014).
 - [5] P. Ade *et al.*, *J. Cosm. Astrop. Phys.* **2019**, 056 (2019).
 - [6] K. Abazajian *et al.*, arXiv e-prints arXiv:1907.04473 (2019).
 - [7] N. Sehgal *et al.*, CMB-HD: An Ultra-Deep, High-Resolution Millimeter-Wave Survey Over Half the Sky, 2019.
 - [8] N. Sehgal *et al.*, CMB-HD: Astro2020 RFI Response, 2020.
 - [9] A. van Engelen *et al.*, *Astrophys. J.* **786**, 13 (2014).
 - [10] S. J. Osborne, D. Hanson, and O. Doré, *J. Cosm. Astrop. Phys.* **2014**, 024 (2014).
 - [11] M. S. Madhavacheril and J. C. Hill, *Phys. Rev. D* **98**, 023534 (2018).
 - [12] S. Ferraro and J. C. Hill, *Phys. Rev. D* **97**, 023512 (2018).
 - [13] E. Schaan and S. Ferraro, *Phys. Rev. Lett.* **122**, 181301 (2019).
 - [14] E. J. Baxter *et al.*, *Phys. Rev. D* **99**, 023508 (2019).
 - [15] O. Darwish *et al.*, arXiv e-prints arXiv:2004.01139 (2020).
 - [16] T. Namikawa, D. Hanson, and R. Takahashi, *Mon. Not. R. Astron. Soc.* **431**, 609 (2013).
 - [17] P. A. R. Ade *et al.*, *Astronomy & Astrophysics* **571**, A17 (2014).
 - [18] W. Hu and T. Okamoto, *Astrophys. J.* **574**, 566 (2002).
 - [19] G. Fabbian, A. Lewis, and D. Beck, *Journal of Cosmology and Astroparticle Physics* **2019**, 057–057 (2019).
 - [20] N. Sehgal *et al.*, *Astrophys. J.* **709**, 920 (2010).
 - [21] S. Foreman, P. D. Meerburg, A. van Engelen, and J. Meyers, *J. Cosm. Astrop. Phys.* **2018**, 046 (2018).
 - [22] J. Dunkley *et al.*, *J. Cosm. Astrop. Phys.* **2013**, 025 (2013).
 - [23] LSST Science Collaboration *et al.*, arXiv e-prints arXiv:0912.0201 (2009).