


Foreground-Immune Cosmic Microwave Background Lensing with Shear-Only Reconstruction

Emmanuel Schaan^{*} and Simone Ferraro[†]

*Lawrence Berkeley National Laboratory, One Cyclotron Road, Berkeley, California 94720, USA
and Berkeley Center for Cosmological Physics, University of California, Berkeley, California 94720, USA*

 (Received 23 April 2018; revised manuscript received 6 August 2018; published 8 May 2019)

Cosmic microwave background (CMB) lensing from current and upcoming wide-field CMB experiments such as AdvACT, SPT-3G and Simons Observatory relies heavily on temperature (versus polarization). In this regime, foreground contamination to the temperature map produces significant lensing biases, which cannot be fully controlled by multifrequency component separation, masking, or bias hardening. In this Letter, we split the standard CMB lensing quadratic estimator into a new set of optimal “multipole” estimators. On large scales, these multipole estimators reduce to the known magnification and shear estimators, and a new shear B -mode estimator. We leverage the different symmetries of the lensed CMB and extragalactic foregrounds to argue that the shear-only estimator should be approximately immune to extragalactic foregrounds. We build a new method to compute, separately and without noise, the primary, secondary, and trispectrum biases to CMB lensing from foreground simulations. Using this method, we demonstrate that the shear estimator is, indeed, insensitive to extragalactic foregrounds, even when applied to a single-frequency temperature map contaminated with cosmic infrared background, thermal Sunyaev-Zel’dovich, kinematic Sunyaev-Zel’dovich, and radio point sources. This dramatic reduction in foreground biases allows us to include higher temperature multipoles than with the standard quadratic estimator, thus, increasing the total lensing signal-to-noise ratio beyond the quadratic estimator. In addition, magnification-only and shear B -mode estimators provide useful diagnostics for potential residuals.

DOI: [10.1103/PhysRevLett.122.181301](https://doi.org/10.1103/PhysRevLett.122.181301)

Introduction.—Weak lensing of the cosmic microwave background (CMB) measures the projected matter distribution throughout the observable Universe and is one of the most promising probes of dark energy, modified gravity, and neutrino masses [1,2]. As the measurement precision increases, systematic biases become more important. While CMB-S4 [3] lensing data should be polarization dominated in the future, in the coming decade, CMB lensing measurements from AdvACT [4], SPT-3G [5], and Simons Observatory [6] will rely heavily on temperature. In this regime, extragalactic foregrounds such as the cosmic infrared background (CIB), the thermal Sunyaev-Zel’dovich effect (TSZ), the kinematic Sunyaev-Zel’dovich effect (KSZ), and radio point sources (PS) can produce biases much larger than the statistical errors, if unaccounted for [7–10]. Mitigation methods have been proposed. For example, masking individually detected or known sources can decrease the bias, and techniques such as bias hardening [8,11,12] are effective when the foreground trispectrum is known. Multifrequency component separation [10] can reduce or null specific foreground components. However, a minimum-variance multifrequency analysis only leads to a modest reduction in foregrounds, and simultaneously nulling TSZ and CIB comes at a large cost, increasing the noise power spectrum by a factor as large as 50 [6]. Furthermore, multifrequency component separation has no effect on the KSZ, which alone causes a significant lensing bias [9]. Therefore, new methods

are needed in order to produce unbiased lensing measurements from CMB temperature maps.

In this Letter, we explore a new approach, leveraging the differing symmetries of the lensing deflections and extragalactic foregrounds in order to separate them. Indeed, as we argue below, extragalactic foregrounds are degenerate with lensing magnification (local monopole distortion of the power spectrum) but not with lensing shear (local quadrupolar distortion) or higher order multipoles. Throughout this Letter, we consider lensing measurements from CMB temperature only, rather than polarization, although we expect a similar approach to work in polarization too.

Lensing multipole estimators.—Estimators: Weak lensing modulates the 2D CMB power spectrum, creating local distortions. These distortions to the power spectrum can be decomposed into a monopole ($m = 0$) corresponding to an isotropic magnification or demagnification, a quadrupole ($m = 2$) corresponding to shearing, as well as higher order even multipoles. Mathematically, the presence of a fixed lensing convergence κ_L , creates off-diagonal correlations in the observed CMB temperature T

$$\langle T_{\ell+\frac{1}{2}} T_{\frac{1}{2}-\ell} \rangle = f_{\ell+\frac{1}{2},\frac{1}{2}-\ell}^{\kappa} \kappa_L + \mathcal{O}(\kappa^2). \quad (1)$$

The angular dependence of the response function f^{κ} can be expanded in multipoles of the angle $\theta_{L,\ell}$ between ℓ and L

$$f_{L+\frac{1}{2},\frac{1}{2}-\ell}^{\kappa} = \sum_{m \text{ even}} f_{L,\ell}^m \cos(m\theta_{L,\ell}), \quad (2)$$




FIG. 1. Noise power spectrum of the lensing convergence κ , reconstructed with the optimal quadratic multipole estimators. Monopole ($m = 0$) and quadrupole ($m = 2$) estimators contain most of the lensing signal-to-noise ratio. The multipole estimators are uncorrelated for $L \lesssim 300$.

which defines the m th multipole response function $f_{L,\ell}^m$. These can be used in Eq. (1) to obtain an estimator of κ_L , from multipole m only. Explicit minimum variance expressions are given in the Supplemental Material [13], and Fig. 1 shows that the monopole and quadrupole estimators contain most of the lensing signal-to-noise ratio, allowing us to neglect estimators with $m > 2$ in practice. To allow a fast evaluation with FFT, we can replace these non-separable optimal multipole estimators by their limits in the “large-scale lens regime”, where large-scale ($L \lesssim 300$) lensing modes are reconstructed from small-scale ($\ell \gtrsim 300$) temperature modes. In this regime, our optimal monopole and quadrupole estimators reduce to the magnification (To be consistent with the optical lensing literature, this estimator should be called “convergence” instead of “magnification”. Since we already use the name convergence to designate the lensing field κ that is being reconstructed, we decided to call shear and magnification the two distinct effects, to avoid confusion.) and shear E -mode estimators of [14–16] (see, also, [14,17–19]), as well as a new shear B -mode estimator

$$\kappa_L = \frac{\int \frac{d^2\ell}{(2\pi)^2} T_\ell T_{L-\ell} g_{L,\ell}}{\frac{2L}{L^2} \cdot \int \frac{d^2\ell}{(2\pi)^2} g_{L,\ell} [\ell C_\ell^0 + (L-\ell) C_{L-\ell}^0]}, \quad (3)$$

where

$$\begin{aligned} g_{L,\ell}^{\text{magnification}} &= \frac{C_\ell^0}{2(C_\ell^{\text{total}})^2} \frac{d \ln \ell^2 C_\ell^0}{d \ln \ell}, \\ g_{L,\ell}^{\text{shear E}} &= \cos(2\theta_{L,\ell}) \frac{C_\ell^0}{2(C_\ell^{\text{total}})^2} \frac{d \ln C_\ell^0}{d \ln \ell}, \\ g_{L,\ell}^{\text{shear B}} &= \sin(2\theta_{L,\ell}) \frac{C_\ell^0}{2(C_\ell^{\text{total}})^2} \frac{d \ln C_\ell^0}{d \ln \ell}. \end{aligned} \quad (4)$$

These estimators should only be interpreted as measuring magnification and shear in the large-scale lens regime ($L \ll \ell$). However, they remain unbiased lensing estimators on all scales. They match the harmonic-space version of [15,16], after normalizing them to be unbiased and with the substitution $T_{\ell+L/2} T_{L/2-\ell} \rightarrow T_\ell T_{L-\ell}$ to allow fast evaluation with FFT. Further, we substitute the lensed CMB power spectrum to C^0 , as is customary for the quadratic estimator (QE) [20–22]. As shown in the Supplemental Material [13] Fig. 1, the magnification and shear estimators are optimal on large scales ($L \lesssim 300$) where they have the same noise as the optimal $m = 0$ and $m = 2$ estimators, are roughly uncorrelated, and recover the signal-to-noise ratio of the standard QE. In the Born approximation, the shear B -mode estimator has zero response to lensing and provides a useful null test. As we show below, it also allows us to detect and subtract any potential “secondary foreground bias” (defined below).

Statistical signal-to-noise ratio: Throughout this Letter, we consider an upcoming stage 3 (“CMB S3”) experiment, with 1.4’ beam FWHM and $7 \mu K$ sensitivity at 148 GHz. We apply the lensing estimators to the single-frequency map at 148 GHz, without any multifrequency component separation. For the lensing weights, we include the lensed CMB, all the foregrounds of the section below titled Sensitivity to foregrounds: simulations, and the detector white noise in the total power spectrum.

Intuitively, Eq. (4) means that magnification can only be measured from a non-scale-invariant power spectrum ($d \ln \ell^2 C_\ell^0 / d \ln \ell \neq 0$), and shear only from a nonwhite power spectrum ($d \ln C_\ell^0 / d \ln \ell \neq 0$). The unlensed CMB power spectrum is neither scale invariant nor white, so a similar signal-to-noise ratio is expected for the shear and magnification estimators. Indeed, as shown in Fig. 2, the

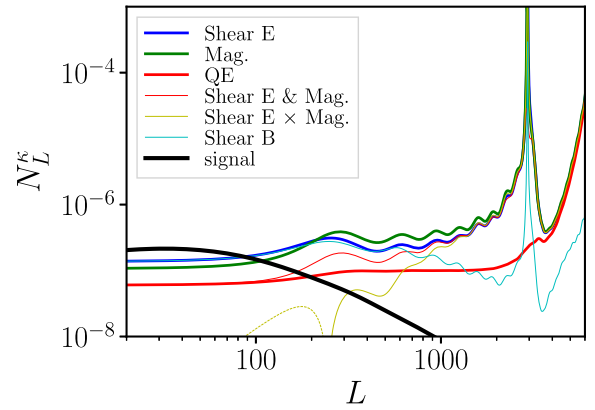


FIG. 2. Lensing reconstruction noise per lensing multipole for the standard quadratic estimator (QE, red), the magnification (green), shear E -mode (blue) and B -mode (cyan) estimators, when using temperature modes $\ell = 30$ –3500. Below multipoles of a few hundred, the shear E and magnification estimators are roughly uncorrelated, and recover the QE when combined, taking into account their noise covariance. Shear E and shear B have similar noise for low multipoles, which makes the shear B a useful null test to compare to shear E .

lensing noise in shear and magnification is comparable. This is convenient: shear and magnification estimators can be compared as a consistency check for residual foregrounds. At fixed $\ell_{\max,T}$, the total signal-to-noise ratio in either shear or magnification is about 60% of that in the QE including the cosmic variance. However, as we show below, the shear estimator is less affected by foregrounds, allowing us to use $\ell_{\max,T} = 3500$ instead of $\ell_{\max,T} = 2500$ for the QE. This allows us to recover all of the signal-to-noise ratio lost by discarding the magnification part. To optimize further, we build a “hybrid estimator” by forming the minimum-variance linear combination of the magnification measured from $\ell_{\max,T} = 2000$ (where foreground contamination is small) and the shear measured from $\ell_T = 30\text{--}3500$. This minimum-variance linear combination takes into account the correlation between the estimators. This “hybrid” estimator, shown in Fig. 3, increases the SNR on the amplitude of lensing by 14% compared to the QE with $\ell_{\max,T} = 2500$, from 93 to 106. A similar hybrid estimator, constructed from the multipole estimators rather than from the magnification and shear will increase the SNR even further.

A spike in the noise power spectrum can be seen for the magnification and shear estimators in Fig. 2, but not for the multipole estimators in Fig. 1. This is a result of the approximate lensing weights in Eq. (4), only valid in the

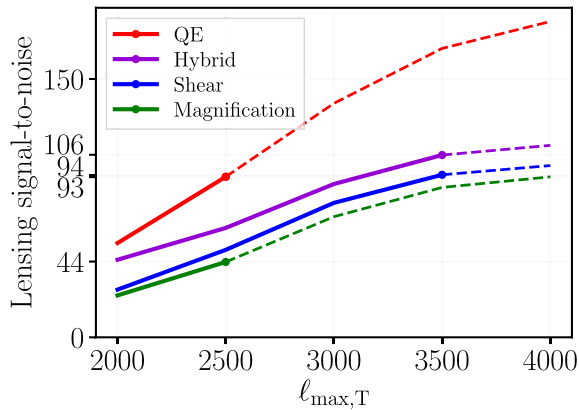


FIG. 3. Total signal-to-noise ratio on the amplitude of the lensing power spectrum, including cosmic variance, as a function of the maximum temperature multipole $\ell_{\max,T}$, for $f_{\text{sky}} = 1$. Different colors correspond to the different estimators. Dashed lines indicate when foreground biases are larger than the statistical uncertainty, even after masking point sources detected at 5σ . At fixed $\ell_{\max,T}$, the signal-to-noise ratio in either shear or magnification is about 60% of the signal-to-noise ratio of the QE. However, as we show below, keeping the foreground bias below the statistical error requires $\ell_{\max,T} = 2500$ for the QE (red dot, $S/N = 70$), compared to $\ell_{\max,T} = 3500$ for the shear estimator (blue dot, $S/N = 77$): hence, the final shear signal-to-noise ratio exceeds that of the QE by 10%. A hybrid estimator QE($\ell \leq 2000$) and shear($\ell = 2000\text{--}3500$) is shown in purple, and increases the signal-to-noise ratio by 37% compared to the standard QE($\ell \leq 2500$).

large-scale lens regime, which cause these estimators to have zero response to lensing (and, thus, infinite noise) at the location of the spike.

Expected sensitivity to foregrounds: Extragalactic foregrounds dominate the lensed CMB on small scales ($\ell \gtrsim 3000$), where they are well described by a one-halo or shot noise term, i.e., by a set of unclustered emission profiles (e.g., halos) or point sources (e.g., galaxies inside azimuthally symmetric halos). If the emission profiles are azimuthally symmetric, the local foreground power spectrum on a small patch of the sky is isotropic, i.e., a function of $\ell = |\ell|$ instead of ℓ . As a result, the corresponding foreground component modifies the observed power spectrum monopole ($m = 0$), but not its higher multipoles. This should bias the magnification estimator and, therefore, the QE, but not the shear estimator.

If the foreground sources are halos with random independent ellipticities, or are pointlike but clustered in elliptical filaments with random orientations, they produce extra noise in the shear estimator, analogous to the shape noise in galaxy lensing. On the other hand, if the ellipticities of foreground halos or of their clustering (filaments) are aligned with the local tidal field, they will produce a bias to the shear estimator, analogously to intrinsic alignments in galaxy lensing (see App. D in [23]).

In summary, any extragalactic foreground biases the magnification estimator and the QE, whereas only foregrounds with specific anisotropies (intrinsic alignments) affect the shear estimators. In the next section, we test this intuition with realistic foreground simulations.

Sensitivity to foregrounds: simulations.—Method: We use simulated maps of lensing convergence, CIB, TSZ, KSZ, and radio PS at 148 GHz from [24], obtained by painting polytropic baryonic profiles on a large-box ($L = 1 \text{ Gpc}/h$) N -body simulation. Crucially, the gas density and temperature profiles given to a halo are not spherical but, instead, follow the triaxiality of the local matter tidal tensor at the position of the halo. As a result, these simulations include a reasonable level of shape noise and “intrinsic alignment”. A halo catalog from this N -body simulation is also available. We reweight these halos to match the redshift distribution of the LSST gold sample, with i -band magnitude $i < 25.3$ [25] [$dn/dz \propto (z/z_0)^2 e^{-z/z_0} / (2z_0)$ with $z_0 = 0.24$], and obtain a projected “galaxy” number density map δ_g . The “galaxy bias” measured from this map roughly matches the expected value $b(z) = 1 + 0.84z$ [25]. These maps have two crucial features: they are realistically correlated with each other, and have a reasonable level of non-Gaussianity. The simulations also include the effect of anisotropic clustering of halos inside filaments, of anisotropic halo profiles, including possible intrinsic alignments. Our goal is to compute the foreground biases to the cross-correlation of CMB lensing with galaxies $C_L^{\kappa\delta_g}$ and to the CMB lensing autospectrum $C_L^{\kappa\kappa}$.

We subtract the mean emission in each foreground map, then rescale the maps by factors of order 1 to match the power spectrum model of [26] (0.38 for CIB, 0.7 for TSZ, 0.82 for KSZ, 1.1 for radio PS). Following [7], we then mask the point sources with flux $\gtrsim 5$ mJy in each foreground map. To do so, we match-filtered the foreground maps with a profile corresponding to the beam and a noise determined by the total power spectrum (lensed CMB plus all foregrounds). The resulting foreground power spectra are shown in the Supplemental Material [13], Fig 3.

In principle, one should add all the foreground maps together to get the total bias, including their correct cross-correlations. However, component separation will reduce each foreground differently. For this reason, we analyze each foreground map separately. This should allow the reader to quantify the foreground bias for any component separation method by rescaling our values appropriately. In what follows, our lens reconstruction relies on temperature multipoles $\ell = 30\text{--}3500$. To measure the lensing bias due to the foregrounds, we decompose the observed sky temperature T_{obs} into the lensed primary CMB T_{CMB} , the foregrounds T_f , and the detector noise T_{noise} : $T_{\text{obs}} = T_{\text{CMB}} + T_f + T_{\text{noise}}$. We write $Q[T_A, T_B]$ for any quadratic estimator (QE, shear, or magnification) applied to maps T_A and T_B , symmetrized in $A \leftrightarrow B$.

As shown in [7–9], biases to the CMB lensing auto power spectrum $C_L^{\kappa\kappa}$ arise from the foreground bispectrum (primary and secondary terms [8]), and from the foreground trispectrum. We evaluate them as follows: (1) The primary bispectrum term is computed as $2\langle Q[T_f, T_f] \kappa_{\text{CMB}} \rangle$, as in [7–9]. (2) The secondary bispectrum could, in principle, be computed as $4\langle Q[T_f, T_{\text{CMB}}] Q[T_f, T_{\text{CMB}}] \rangle$. However, this autocorrelation is biased by the large noise of $Q[T_f, T_{\text{CMB}}]$, which would have to be subtracted accurately. Therefore, we propose and implement a new method to avoid this issue. We Taylor-expand the lensed CMB map $T_{\text{CMB}} = T^0 + T^1 + \dots$ in powers of κ and compute the quantity $8\langle Q[T_f, T^0] Q[T_f, T^1] \rangle$ (Another way to evaluate the secondary bispectrum term would be $\langle Q[T_f, T_{\text{CMB}}] Q[T_f, T_{\text{CMB}}] - Q[T_f, T'_{\text{CMB}}] Q[T_f, T'_{\text{CMB}}] \rangle$ where T'_{CMB} is constructed from the same unlensed CMB realization as T_{CMB} but lensed by an independent κ realization.). This works because the quadratic estimators are, by construction, unbiased when applied to the pair (T^0, T^1) , to first order in lensing. This greatly reduces the noise, and this is a cross-correlation so no noise subtraction is needed (no N^0 , or higher order bias N^i). (3) For the trispectrum term, we compute $\langle Q[T_f, T_f] Q[T_f, T_f] \rangle$, and subtract the Gaussian contribution (which is a part of N^0) analytically, as in [7,8].

For the cross-correlation with tracers $C_L^{\kappa\delta_g}$, only the primary bispectrum is present, and without the combinatorial factor 2: $\langle Q[T_f, T_f] \delta_g \rangle$. The secondary bispectrum and trispectrum terms only act as a source of noise on this cross-correlation, not bias.

Results: The resulting foreground biases for the cross-correlation $C_L^{\kappa\delta_g}$ are shown in Fig. 4. Despite the masking of point sources, the CIB, TSZ, KSZ, and radio PS lead to very large and statistically significant biases for the QE and the magnification estimators. Again, multifrequency component separation may be used to null the TSZ bias, or reduce the CIB or radio PS biases. However, reducing all these biases simultaneously typically causes a large noise increase. Furthermore, multifrequency analyses have no effect on the KSZ bias. Therefore, these foreground biases are a major concern for the standard QE. On the other hand, no foreground bias is detected in the shear estimator. This is the main result of this Letter: even when applied to a single-frequency temperature map, the shear estimator measures only the quadrupolar distortions from lensing and, therefore, is immune to foregrounds. It is remarkable that this holds even for a single frequency map out to $\ell_{\text{max},T} = 3500$, where the temperature modes are foreground dominated. Our QE TSZ bias in Fig. 4 is smaller than in [10,27], which can be explained by our scaling down of the TSZ map to match the power spectrum model of [26], our masking, and the different redshift of our galaxy catalog. Our CIB bias is slightly larger than found in [27].

For the lensing autospectrum $C_L^{\kappa\kappa}$, the primary, secondary, and trispectrum biases discussed in the previous section are shown in Fig. 5. At low (respectively. high)




FIG. 4. Relative bias to the cross-correlation between CMB lensing and the LSST gold galaxy sample, as a function of lensing multipole L , when including temperature multipoles $\ell = 30\text{--}3500$ at 148 GHz. This foreground bias corresponds to the primary bispectrum term. The grey boxes indicate bins of lensing multipoles with the corresponding statistical error bars for the standard quadratic estimator (lensing noise plus cosmic variance, identical in each panel). The foreground biases are much larger than the statistical error bars for both the standard quadratic estimator and the magnification estimator, whereas they are barely measurable for the shear estimator.




FIG. 5. Relative foreground bias on the CMB lensing power spectrum, as a function of lensing multipole L , when including temperature multipoles $\ell = 30\text{--}3500$ at 148 GHz. The grey boxes indicate bins of lensing multipoles with the corresponding statistical error bars for the standard quadratic estimator (lensing noise plus cosmic variance). Top: primary bispectrum bias, dominant at low L . Middle: trispectrum bias, dominant at high L . Bottom: secondary bispectrum bias. The dominant biases (primary and trispectrum) are much larger than the statistical error bars for the QE and magnification estimator and are barely measurable for the shear estimator. The secondary bispectrum bias is smaller and similar in size for all estimators. The secondary bispectrum bias is identical for the shear E and B estimators, making the difference of the two an unbiased lensing estimator.

lensing multipoles, the primary (respectively. trispectrum) bias dominates. In both cases, a large bias is seen in the QE and magnification estimator, while the shear estimator is unbiased. Our primary and trispectrum foreground biases are consistent with the results of [7] for the CIB and TSZ, and slightly smaller than what is found in [9] for the KSZ, due to our rescaling of the KSZ map and the slightly different lensing weights. We compute the secondary foreground bias separately. This term is smaller than the primary and trispectrum term, but non-negligible for L of a few hundred. Here, the shear estimator alone does not improve over the QE and magnification estimators. This occurs because the shear secondary bias introduces a $\cos^2(2\theta)$, which makes it sensitive to the foreground monopole power. However, the shear B -mode estimator has the same secondary bias and no response to lensing: therefore, subtracting it from the shear E mode cancels the secondary bias at the cost of an increased noise. Overall, the shear estimator dramatically reduces the foreground biases. In the absence of any foreground cleaning, the shear estimator allows us to increase the range of multipoles used in the lens reconstruction from $\ell_{\max,T} \approx 2500$ for the QE, to $\ell_{\max,T} \approx 3500$ for shear only. Multifrequency foreground cleaning may help increase the range of usable multipoles—and, thus, the statistical power—for all estimators. The proposed shear B -mode subtraction may further improve the range for the shear E -mode estimator. We leave a detailed optimization study to future work.

Conclusion.—For current and upcoming CMB experiments such as AdvACT, SPT-3G, and Simons Observatory, CMB lensing reconstruction will rely heavily on temperature. Foreground emission is known to contaminate temperature maps from which lensing is reconstructed and, therefore, produce very significant biases, leading to wrong conclusions about cosmology if unaccounted for. Modeling and subtracting these bias terms is likely to be very challenging, due to the complex baryon physics involved in producing them. While some foregrounds can be nulled (TSZ) or reduced (CIB, radio PS) by a multifrequency analysis at the cost of a degradation in map noise, other foregrounds cannot (KSZ).

Therefore, in this Letter, we explored a different approach, by using the approximate isotropy of the extragalactic foreground 2D power spectra, and splitting the QE into optimal quadratic multipole estimators.

In the large-scale lens regime, they reduce to the isotropic magnification and anisotropic shear E -mode estimators of [14–16], and a new shear B -mode estimator. The shear estimator enables a remarkable reduction of foreground biases, compared to the QE, even when applied to a single-frequency temperature map. As a result, the shear estimator allows us to increase the range of multipoles used in the lens reconstruction to $\ell_{\max,T} \approx 3500$ instead of $\ell_{\max,T} \approx 2500$ for the QE, while keeping foreground biases within the statistical uncertainty. Overall, the

signal-to-noise ratio in shear with $\ell_{\max,T} = 3500$ is very similar to that in QE with $\ell_{\max,T} = 2500$. Thus, the shear estimator provides a robust way of measuring lensing. Component separation may allow the use of higher multipoles for all estimators. On the other hand, the magnification estimator is highly sensitive to foregrounds, so comparing magnification and shear provides an excellent diagnostic for foreground contamination. The shear B -mode estimator constitutes an additional null test, and allow us to further reduce foreground biases. Quantifying the size of the higher order biases such as $N^{(1)}$ and $N^{(2)}$ for the shear and magnification estimators will be important.

Further optimization is possible, by combining different estimators with different $\ell_{\max,T}$. For instance, a hybrid estimator magnification ($\ell \leq 2000$) and shear ($\ell \leq 3500$) improves the lensing signal-to-noise ratio by 14% compared to the standard QE ($\ell \leq 2500$).

Better approximations to the optimal multipole estimators than the shear and magnification estimators may yield further improvements in signal-to-noise ratio. A promising approach would be to replace the derivatives in Eq. (4) by free functions of ℓ to be optimized. Future CMB lensing data from CMB S4 should be polarization dominated. The shear and magnification estimators can be generalized to polarization [16], and may bring improvements there too. This would have implications for precision delensing, in order to isolate primordial tensor modes. Similar foreground biases occur in lens reconstruction from intensity mapping [23,28] (e.g., the “self-lensing bias” for CIB), and the shear estimator may allow us to reduce them [23,28]. Finally, the split into magnification and shear E and B modes may also help detect residual Galactic foregrounds or beam ellipticity. We leave the exploration of these promising avenues to future work.

We thank Marcelo Alvarez, Anthony Challinor, Sandrine Codis, Simon Foreman, Colin Hill, Shirley Ho, Akito Kusaka, Antony Lewis, Heather Prince, Uroš Seljak, David Spergel, Blake Sherwin, Alex van Engelen, Martin White, and Hong-Ming Zhu for useful discussion. We thank the anonymous referees for very useful comments and suggestions, which greatly improved this Letter. E.S. is supported by the Chamberlain fellowship at Lawrence Berkeley National Laboratory. S.F. was in part supported by a Miller Fellowship at the University of California, Berkeley and by the Physics Division at 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.

*eschaan@lbl.gov

†sferraro@lbl.gov

- [1] A. Lewis and A. Challinor, *Phys. Rep.* **429**, 1 (2006).
- [2] D. Hanson, A. Challinor, and A. Lewis, *Gen. Relativ. Gravit.* **42**, 2197 (2010).
- [3] K.N. Abazajian, P. Adshead, Z. Ahmed *et al.*, [arXiv:1610.02743](https://arxiv.org/abs/1610.02743).
- [4] S.W. Henderson, R. Allison, J. Austermann *et al.*, *J. Low Temp. Phys.* **184**, 772 (2016).
- [5] B. A. Benson, P. A. R. Ade, Z. Ahmed *et al.*, *Proc. SPIE Int. Soc. Opt. Eng.* **9153**, 91531P (2014).
- [6] P. Ade *et al.* (Simons Observatory Collaboration), *J. Cosmol. Astropart. Phys.* **02** (2019) 056.
- [7] A. van Engelen, S. Bhattacharya, N. Sehgal, G. P. Holder, O. Zahn, and D. Nagai, *Astrophys. J.* **786**, 13 (2014).
- [8] S.J. Osborne, D. Hanson, and O. Doré, *J. Cosmol. Astropart. Phys.* **03** (2014) 024.
- [9] S. Ferraro and J.C. Hill, *Phys. Rev. D* **97**, 023512 (2018).
- [10] M.S. Madhavacheril and J.C. Hill, *Phys. Rev. D* **98**, 023534 (2018).
- [11] T. Namikawa, D. Hanson, and R. Takahashi, *Mon. Not. R. Astron. Soc.* **431**, 609 (2013).
- [12] S. Foreman, P.D. Meerburg, A. van Engelen, and J. Meyers, [arXiv:1803.04975](https://arxiv.org/abs/1803.04975).
- [13] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevLett.122.181301> for the full derivation, implementation and tests of the lensing estimators.
- [14] T. Lu and U.-L. Pen, *Mon. Not. R. Astron. Soc.* **388**, 1819 (2008).
- [15] M. Bucher, C.S. Carvalho, K. Moodley, and M. Remazeilles, *Phys. Rev. D* **85**, 043016 (2012).
- [16] H. Prince, K. Moodley, J. Ridl, and M. Bucher, *J. Cosmol. Astropart. Phys.* **01** (2018) 034.
- [17] M. Zaldarriaga and U. Seljak, *Phys. Rev. D* **59**, 123507 (1999).
- [18] U.-L. Pen, *New A* **9**, 417 (2004).
- [19] T. Lu, U.-L. Pen, and O. Doré, *Phys. Rev. D* **81**, 123015 (2010).
- [20] D. Hanson, A. Challinor, G. Efstathiou, and P. Bielewicz, *Phys. Rev. D* **83**, 043005 (2011).
- [21] A. Lewis, A. Challinor, and D. Hanson, *J. Cosmol. Astropart. Phys.* **11** (2011) 18.
- [22] W. Hu and T. Okamoto, *Astrophys. J.* **574**, 566 (2002).
- [23] S. Foreman, P.D. Meerburg, A. van Engelen, and J. Meyers, *J. Cosmol. Astropart. Phys.* **07** (2018) 046.
- [24] N. Sehgal, P. Bode, S. Das, C. Hernandez-Monteagudo, K. Huffenberger, Y.-T. Lin, J.P. Ostriker, and H. Trac, *Astrophys. J.* **709**, 920 (2010).
- [25] P. A. Abell, J. Allison *et al.*, [arXiv:0912.0201](https://arxiv.org/abs/0912.0201).
- [26] J. Dunkley, E. Calabrese, J. Sievers *et al.*, *J. Cosmol. Astropart. Phys.* **07** (2013) 025.
- [27] E. J. Baxter, Y. Omori, C. Chang *et al.*, *Phys. Rev. D* **99**, 023508 (2019).
- [28] E. Schaan, S. Ferraro, and D. N. Spergel, *Phys. Rev. D* **97**, 123539 (2018).