

CMB lensing reconstruction biases from masking extragalactic sourcesMargherita Lembo^{1,2,3,*} Giulio Fabbian^{4,5,1} Julien Carron^{6,1} and Antony Lewis¹¹*Department of Physics and Astronomy, University of Sussex, Brighton BN1 9QH, United Kingdom*²*Dipartimento di Fisica e Scienze della Terra, Università degli Studi di Ferrara,
via Giuseppe Saragat 1, I-44122 Ferrara, Italy*³*Istituto Nazionale di Fisica Nucleare (INFN), Sezione di Ferrara,
Via Giuseppe Saragat 1, I-44122 Ferrara, Italy*⁴*Center for Computational Astrophysics, Flatiron Institute, 162 5th Avenue, New York, New York 10010 USA*⁵*School of Physics and Astronomy, Cardiff University, The Parade, Cardiff CF24 3AA, United Kingdom*⁶*Université de Genève, Département de Physique Théorique et CAP,
24 Quai Ansermet, CH-1211 Genève 4, Switzerland*

(Received 28 September 2021; accepted 13 April 2022; published 26 July 2022)

Observed cosmic microwave background (CMB) temperature and polarization maps can be powerful cosmological probes and used for CMB lensing reconstruction. However, the CMB maps are inevitably contaminated by foregrounds, some of which are usually masked to perform the analysis. If this mask is correlated to the lensing signal, measurements over the unmasked sky may give biased estimates and hence biased cosmological inferences. For example, masking extragalactic astrophysical emissions associated with objects located in dark matter halos will systematically remove parts of the sky that have a mass density higher than average. This can lead to a modified lensed CMB power spectrum over the unmasked area and biased measurements of lensing reconstruction auto- and cross-correlation power spectra with external matter tracers (both from the direct impact on the lensing power and via modifications to the lensing-dependent reconstruction power spectrum corrections, $N_L^{(0)}$, $N_L^{(1)}$, and $N_L^{(3/2)}$). For direct masking of the CMB lensing field, we give an approximate halo-model prediction of the size of the effect and derive simple analytic models for point sources and threshold masks constructed on a correlated Gaussian foreground field. We show that biases are significantly reduced by optimal filtering of the CMB maps in the lensing reconstruction, which effectively fills back some of the information in small mask holes. We test the resulting lensing power spectrum biases on numerical simulations, masking radio sources, and peaks of thermal Sunyaev-Zeldovich (tSZ) and cosmic infrared background (CIB) emission. For radio point sources, the remaining bias is negligible, but a temperature lensing reconstruction power spectrum bias remains at the 0.5%–5% level if clusters with a mass larger than $\sim 10^{14} M_\odot/h$ are masked or if peaks in the CIB and tSZ maps are removed with $f_{\text{sky}}^{\text{mask}} \simeq 0.5\text{--}7\%$. In any case, these biases can only be measured with a statistical significance $\lesssim 2\sigma$ for future datasets. Moreover, we quantified the impact of the mask biases in the cross-correlation power spectrum between CMB lensing and tSZ and CIB and found them to be larger (up to $\sim 30\%$). We found that masking tSZ-selected galaxy clusters leads to the largest mask biases, potentially detectable with high significance, and should therefore be avoided as much as possible. For the most realistic masks we considered, masking biases can only be measured with marginal significance. We found that the calibration of cluster masses using CMB lensing, in particular for objects at $z \lesssim 0.6$, might be significantly affected by mask biases for near-future observations if the lensing signal recovered inside the mask holes is used without further corrections. Conversely, mass calibration of high redshift objects will still deliver unbiased results.

DOI: [10.1103/PhysRevD.106.023525](https://doi.org/10.1103/PhysRevD.106.023525)**I. INTRODUCTION**

Cosmic microwave background (CMB) observations are inevitably contaminated at some level by foregrounds, from galactic dust to a range of extragalactic signals

including the cosmic infrared background (CIB), the Sunyaev-Zeldovich effect (SZ), and radio point sources. Extragalactic emissions are correlated to the large-scale structure from which they originate and hence correlate with other tracers of the matter distribution over a similar redshift range [1–11]. Gravitational lensing is an important effect on the CMB. It both modifies the observed power

*M.Lembo@sussex.ac.uk

spectra and produces non-Gaussian signatures in the CMB maps that can be used for lensing reconstruction. The lensing signal correlates with the extragalactic foregrounds, and this means that foreground masking has the potential to produce biased inferences from observations of the lensed CMB over the unmasked area.

In principle, the nonblackbody spectrum of the foregrounds (other than kinetic SZ) can be used to clean the foregrounds through multifrequency observations. Foreground cleaning has been very successfully used, particularly on large scales, but inevitably comes with the cost of increased noise, especially on small scales where the observational noise becomes comparable to the observed signal. Alternatively, the foreground signal can simply be modeled, which is what is often done at the power spectrum level for CMB likelihood analysis. In both cases, it is often necessary to also apply some masking to the brightest sources, including the Galactic plane (which is not correlated to large-scale structure and hence does not introduce a direct bias) but also extragalactic sources. For CMB lensing reconstruction, the non-Gaussianity of the foregrounds is important and can produce a direct bias on lensing estimates that is difficult to model [12,13]. For thermal Sunyaev-Zeldovich (tSZ) foregrounds, the largest non-Gaussianity is associated with the brightest tSZ clusters and hence can be substantially reduced by cluster masking [12,13]. Point sources and CIB foreground non-Gaussianity can also be substantially reduced by masking the brightest sources or through component separation. A variety of other methods have been suggested to reduce lensing biases from small-scale temperature reconstruction [14–18]; however, these are usually only applied after the brightest sources have already been masked out. To extract reliable information from small-scale CMB temperature observations, it is therefore likely to be necessary to understand the impact of the source masking, especially if the correlated masking introduces substantial biases.

Recent Atacama Cosmology Telescope (ACT) lensing analyses successfully performed lensing reconstruction on maps where extragalactic sources were subtracted from the map with a dedicated procedure, but not masking the corresponding sky areas [16,19]. Despite the success of the method, it is unclear if it will remain sufficiently accurate for the analysis of future high sensitivity experiments. Source subtraction may itself alter the underlying CMB signal on small scales and become problematic, in particular for extended tSZ-detected clusters, and dedicated assessment of potential systematics introduced by the technique would have to be carried out anyway.

Masking the observed CMB data with a lensing-correlated mask introduces a number of different effects. First, as shown in Ref. [20] (hereafter, Paper I), the CMB power spectra on the unmasked area can be significantly altered on small scales (and large scales for B-mode polarization), since there is a net scale-dependent

demagnification over the unmasked area. Although estimated to be negligible for *Planck* [21], the impact is potentially larger for forthcoming high-resolution CMB experiments where more of the information is in the small scales, precisely where the foregrounds are more important. This has the potential to bias parameter constraints if not consistently accounted for. Second, if the CMB lensing potential is estimated using standard quadratic estimators [22], the mask complicates its normalization and the noise biases to the lensing power spectrum, even for masks that are uncorrelated to the signal. For uncorrelated masks, these effects can be estimated analytically in some cases (see Appendix A) or quantified for a fiducial model using independent lensed CMB simulations. After correcting for these effects, if the mask is actually correlated to the signal, there will be additional biases because (1) the areas of high convergence have been preferentially masked directly biasing the actual lensing power over the unmasked area; (2) the normalization, mean-field, and Gaussian noise bias ($N_L^{(0)}$) are different because of the differences in the lensed CMB power spectra over the unmasked areas; (3) the $N_L^{(1)}$ bias due to nonprimary lensing contractions [23] is changed due to the different CMB and lensing power; (4) non-Gaussian biases, specifically $N_L^{(3/2)}$ related to nonlinear large-scale structure growth and post-Born lensing effects [18,24–27], are modified due to the changed non-Gaussianity (e.g., reduced skewness) of the masked field and related changes in the power spectra.

Effects of large-scale structure (LSS)-correlated masking in the lensing reconstruction might also impact the estimated CMB lensing field itself and therefore the cosmological constraints involving its statistics beyond the power spectrum [28]. Moreover, it might bias the cross-correlation with external matter tracers. These are powerful cosmological and astrophysical probes for a wide range of physical phenomena. Cross-correlation between CMB lensing and galaxy surveys data helps improve cosmological constraints by breaking parameter degeneracies (e.g., involving galaxy bias) and by measuring nuisance parameters associated with sources of systematic errors (e.g., lensing multiplicative bias and photometric redshift errors) [29–32]. For future galaxy surveys, such as Euclid and LSST, this approach is likely to become the standard analysis to obtain cosmological constraints [33,34].

The cross-correlation between CMB lensing and the extragalactic emissions can also be a useful source of additional information. The tSZ-CMB lensing cross-correlation ($\kappa \times y$) signal is a unique probe of the physics of the intracluster medium at high-redshift $z \approx 1$ and in relatively low-mass clusters and groups ($10^{13} \lesssim M \lesssim 10^{15} M_\odot/h$). It is also more sensitive to contributions from structures located in dark matter halos of lower masses than the tSZ autopower spectrum or the cross-correlation with galaxy lensing measurements [35–38].

Furthermore, it is a powerful cosmological probe due to its strong dependency on σ_8 and Ω_m [39].

The CIB emission, conversely, is generated by redshifted thermal radiation from UV-heated dusty star-forming galaxies (DSFGs) that have a redshift distribution peaked between $1 \lesssim z \lesssim 2$ [40,41]. The kernel of CMB lensing peaks around the same redshift range, and CMB photons are mainly lensed by halos of mass $\approx 10^{11} - 10^{13} M_\odot$ similar to the one of DSFGs [11]. As such, CIB and CMB lensing are highly correlated. Their cross-correlation ($\kappa \times$ CIB) was first measured by *Planck* [42] in multiple frequency bands with high statistical significance. Such strong correlation, as high as 80% at 545 GHz, allows the star formation rate density and the mass of the halos hosting CIB to be constrained over a wide range of redshifts, $1 \lesssim z \lesssim 7$. Future CMB lensing and CIB measurements of Simons Observatory (SO) [43] and CCAT-prime [44] are expected to further improve the constraining power on these astrophysical processes [45]. The high degree of correlation between CIB and CMB lensing can also be exploited to construct templates of the lensed B-modes for delensing analyses using CIB as an external LSS tracer alone or in combination with CMB lensing itself [46,47]. Potential biases due to LSS-correlated masking on $\kappa \times$ CIB would not directly translate in a misinterpretation of the delensed B-modes signal (and hence in a bias on r) as the correlation between CIB and CMB lensing is usually not assumed from a model but directly measured from the data themselves. Nevertheless, it is important to understand if these effects might impact the expected delensing efficiency for future experiments.

Building on our work of Paper I, in this paper, we study and quantify the impact of the mask bias for lensing reconstruction. The main aim of this paper is to identify the important sources of correlated mask bias, provided some qualitative understanding of them, and approximately quantify their size. Future work will be required to make accurate and fully self-consistent predictions including the impact of potential foreground residuals in the masked CMB in addition to the effect of the lensing-correlated mask.

In Sec. II, we give simple analytic models of the expected effect of foreground masking for various types of mask, taking the masks to apply directly to the lensing convergence field without modeling lensing reconstruction. In Sec. III, we describe the realistic simulations, including extragalactic foreground emissions as well as the effect of nonlinear evolution of LSS, that we used to model the relevant effects and validate our analytic models. A reader who is only interested in the actual results on lensing reconstruction can skip these two first sections and go directly to Sec. IV, where we show the impact of the mask bias on the final estimated lensing potential power spectrum from lensed CMB maps, including the effects of optimal CMB filtering and lensing reconstruction.

In Sec. V, we also quantify the impact on the cross-correlation power spectrum with the true convergence field and extragalactic foregrounds and assess the impact on cluster mass calibration if the lensing field recovered inside a cluster mask is used directly without further modeling. Finally, in Sec. VI, we estimate the impact of these biases on the near-future CMB measurements. Modeling the full effect of mask bias on the reconstructed CMB lensing potential field is challenging analytically, but we give some analytic results in Appendix A for the case of uncorrelated circular mask holes.

For near-future observations, such as SO, the temperature signal still contains a substantial fraction of the available information, so fully exploiting the data will require robust modeling of the temperature signal, which is what we focus on in this work. The contamination induced by foregrounds is less important for CMB lensing reconstruction using polarization [48], since the polarized foreground amplitudes are expected to be substantially lower, but should also be assessed in future work [future observations by CMB-S4 (S4 hereafter) [49] will give lensing reconstructions that are more dominated by polarization].

II. LENSING POWER SPECTRA WITH LSS-CORRELATED MASKS

A simple estimate of the effect of an LSS-correlated mask on the CMB lensing power can be obtained by estimating the power spectrum of the true κ field over the unmasked area and comparing its value with the one computed over the full sky. This operation can only be performed on simulations, since κ is not measurable directly but has to be reconstructed from lensed CMB maps. However, direct κ masking is more easily modeled analytically, though it would only provide a good estimate of the expected amplitude of the effect if the lensing reconstruction were a fully local noiseless unbiased estimate of the true underlying field. We use it as a simple model to gain some analytic insight into the impact of correlated masking without the complications induced by the lensing reconstruction estimator. We give simple analytic models of the direct masking effects for all the masks described in the next section and also compare these analytic predictions with measurements based on direct masking of simulations. In the following, we assume sky masks to be uncorrelated with the unlensed CMB and ignore the correlation between CMB lensing and CMB temperature (i.e., $C_L^{T\phi} = 0$) induced by the integrated Sachs-Wolfe effect in both the analytical modeling and in the simulation measurements.

A. Halo model

A qualitative explanation of the effect of masking peaks of the density field can be obtained using the halo model of

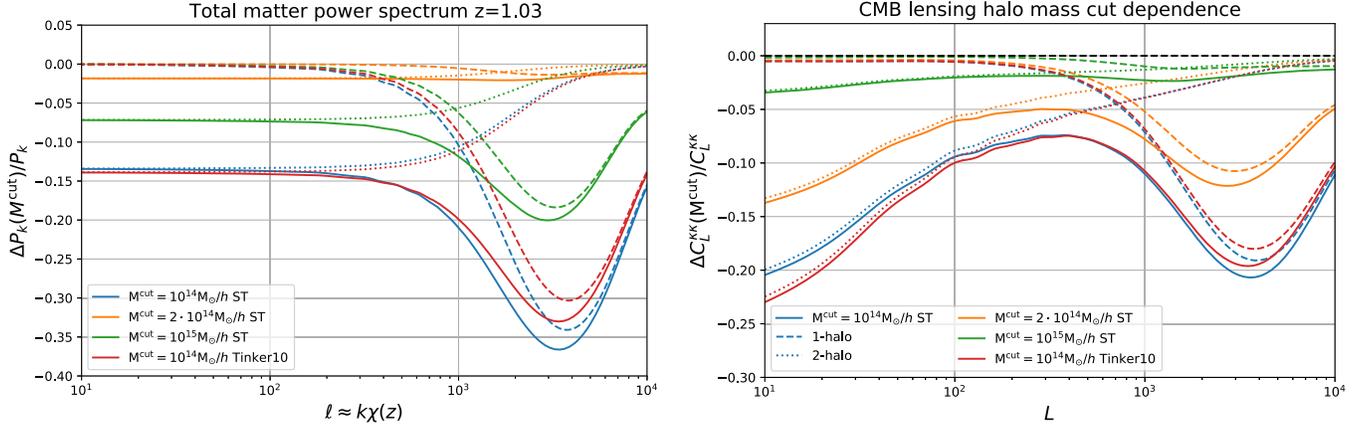


FIG. 1. Left: fractional change in the halo-model matter power spectrum due to cutting the most massive haloes at $z \approx 1$, close to the peak of the CMB lensing kernel. The ratio is shown as function of $k\chi \sim \ell$. Right: fractional impact of the most massive halos on the corresponding CMB lensing convergence angular power spectrum. In both panels, the dashed and dotted lines show the $1h$ and $2h$ contributions, which sum up to the solid lines. We show results for both Sheth and Tormen [53,56,57] (ST) and Tinker *et al.* 2010 [58] mass functions. The choice of the mass function has little impact on the results.

large structure (see, e.g., Ref. [50] for a review). In this model, the entire matter of the Universe is distributed inside haloes of roughly universal spherically averaged halo profile $\rho_M(r)$ as a function of the halo mass. In the following, we assume $\rho_M(r)$ to be a Navarro-Frenk-White profile [51] defined by the mass-concentration relation of Ref. [52]. The matter power spectrum $P(k)$ is built as the sum of two separate parts, the one- and two-halo terms ($1h$ and $2h$, respectively), that correlate points within the same or two different halos respectively. This reads

$$P(k) \equiv P_{1h}(k) + P_{2h}(k) \quad (2.1)$$

$$P_{1h}(k) \propto \int dM \frac{dn}{dM} |\rho_M(k)|^2 \quad (2.2)$$

$$P_{2h}(k) \propto P^{\text{lin}}(k) \left(\int dM \frac{dn}{dM} b(M) \rho_M(k) \right)^2, \quad (2.3)$$

where M is the halo mass, dn/dM is the halo mass function which describes the number density of halos of a given mass, $b(M)$ is the halo bias, $\rho_M(k)$ is the Fourier transform of $\rho_M(r)$, and $P^{\text{lin}}(k)$ is the linear matter power spectrum. To compute the integrals of Eqs. (2.2) and (2.3), we used the public HMVEC code¹ and a Sheth and Tormen [53] mass function. To emulate the masking of a nearly mass-limited sample of halos, we truncate the integration over the halo mass function in both terms of Eq. (2.1) to different mass limits, such as those expected for the SO and S4 tSZ-selected clusters (see the next section for more details). This approximation can be refined to include selection function effects of each experiment [54]. The signature of such a cut

is twofold: a large direct suppression of power on small scales from the $1h$ term, and, since the most massive halos form in the highest peaks of the linear density and are highly biased, a suppression on large scales as well.

In Fig. 1, we show the ratio between the two-dimensional projection of the density power spectra in the Limber approximation after the truncation in the integration, and our reference power spectrum that (very conservatively) uses $10^{18} M_\odot$ as the upper limit of the integration for two different mass cuts (blue and orange). The solid lines show the ratio of the spectra including both the $1h$ and $2h$ terms at $z = 1.03$; the dashed lines show the effect of the mass cut on $1h$ term, while the dotted lines show the same effect for the $2h$ term. The suppression of power on small scales is visible, as is a large dip at $\ell \simeq k\chi \sim 3000$, while the $2h$ contribution dominates on large scales. Since our goal is to assess the impact of masking on the CMB lensing, we therefore used the matter power spectrum with and without the mass cut in the computation of the CMB lensing convergence power spectrum

$$C_L^{\kappa\kappa} = \frac{9H_0^4 \Omega_{m,0}^2}{4c^4} \int_0^{\chi_s} d\chi \left(\frac{\chi_s - \chi}{\chi_s} \right)^2 P \left(\frac{L + 1/2}{\chi}, \chi \right), \quad (2.4)$$

where χ is the comoving distance and χ_s is the comoving distance to the CMB. As noted in Ref. [55], the $2h$ contribution is not naturally built to match the expectation of linear cosmological perturbation theory. We follow their approach and normalize the $2h$ contribution to match the linear theory prediction when computed over our reference range of integration and apply consistently the same normalization to the cases when the integration range is cut. We regularized the $1h$ term on large scales by exponentially damping the contribution of matter

¹<https://github.com/simonsobs/hmvec>.

perturbations on scales $k \leq 0.01$ Mpc/ h in the integrand of Eq. (2.4) [50]. The overall change observed in $C_L^{\kappa\kappa}$ is a consequence of the superposition of the power suppression observed for the matter power spectrum at different redshifts weighted by the lensing kernel.

It may also be possible to extend the halo model approach to make predictions for masking bright infrared sources, or a more direct model of the tSZ intensity. However, we do not pursue this further here, since, as we discuss in Sec. IV, the actual response of the lensing reconstruction to masking is rather different due to the filtering on the CMB maps [41].

B. Gaussian foreground peaks model

When the mask is obtained by masking the peaks of a foreground field, we can construct an analytic model by approximating both the foreground field and the lensing convergence as being Gaussian, with some known covariance. More generally, we can consider the effect of masking on the cross-correlation function of any two Gaussian fields (which we could take to both be the lensing convergence or, for a cross-correlation case, the lensing convergence and some tracer of the matter density). In real space, following a similar argument to Paper I, the main quantity of interest is the masked correlation function of two fields A and B ,

$$\xi_{\text{masked}}^{AB}(r) \equiv \langle W(\mathbf{x})A(\mathbf{x})W(\mathbf{x}')B(\mathbf{x}') \rangle, \quad (2.5)$$

where $\mathbf{x}' = \mathbf{x} + \mathbf{r}$ and $W(\mathbf{x})$ is a mask window function. We assume that some underlying Gaussian statistically isotropic “foreground”: field $f(\mathbf{x})$ (i.e., the κ , y and CIB fields) determines the mask probability locally, so that $W(\mathbf{x})$ only depends on some (in general nonlinear) function of $f(\mathbf{x})$ at the same point. For a mask that is constructed by thresholding the foreground to remove the peaks of its emission, the mask window function is given by a simple step function $W(\mathbf{x}) = \Theta(\nu\sigma_f - f(\mathbf{x}))$ where ν determines the “sigma” value of the cut. The expectation value in Eq. (2.5) is then an integral over the correlated Gaussian variables $f(\mathbf{x})$, $f(\mathbf{x}')$, $A(\mathbf{x})$, and $B(\mathbf{x}')$, where the components of the covariance matrix are simply the correlation functions of the fields evaluated at \mathbf{r} , or zero. If we know these correlation functions, the expectation value can therefore be calculated straightforwardly as the two Gaussian integrals over $A(\mathbf{x})$ and $B(\mathbf{x}')$ can be done directly.

For Gaussian fields with any local mask (in the sense defined above), the bias may be written compactly by introducing the Gaussian independent variable $f_{\pm} = f(\mathbf{x}) \pm f(\mathbf{x}')$ and integrating analytically over the correlated Gaussian distributions. The bias to the mask-deconvolved correlation function, an estimate of the full-sky correlation function, is then

$$\begin{aligned} \Delta_{\xi}^{\xi^{AB}}(r) &\equiv \frac{\xi_{\text{masked}}^{AB}(r)}{\langle W(\mathbf{x})W(\mathbf{x}') \rangle} - \xi^{AB}(r) \\ &= \frac{\xi^{Af_+}(r)\xi^{Bf_+}(r) \langle (f_+^2 - \sigma_+^2)W(\mathbf{x})W(\mathbf{x}') \rangle}{\sigma_+^4(r) \langle W(\mathbf{x})W(\mathbf{x}') \rangle} \\ &\quad - \frac{\xi^{Af_-}(r)\xi^{Bf_-}(r) \langle (f_-^2 - \sigma_-^2)W(\mathbf{x})W(\mathbf{x}') \rangle}{\sigma_-^4(r) \langle W(\mathbf{x})W(\mathbf{x}') \rangle}. \end{aligned} \quad (2.6)$$

In this equation, $\xi^{Xf_{\pm}}(r)$ is $\xi^{Xf}(0) \pm \xi^{Xf}(r)$, and $\sigma_{\pm}^2(r) = 2\sigma_f^2 \pm 2\xi_f(r)$. For a specific given mask, Eq. (2.6) provides an empirical estimate of the bias using simulations or data, after computation of spectra and cross-spectra of $W(\mathbf{x})$, $(fW)(\mathbf{x})$, and $(f^2W)(\mathbf{x})$. In the case of a threshold mask, f_- is left unconstrained from the threshold windowing, and f_+ is restricted not to exceed $2\nu\sigma_f$. The averages required result in one-dimensional, very smooth integrals which are easily computed numerically, giving us an alternative fully analytic prediction.

In the limit of large separations, where the foreground fields at the two points are uncorrelated, taking $\sigma_f^2 \gg \xi_f(r)$, $|\xi^{Af}(r)| \ll |\xi^{Af}(0)|$, $|\xi^{Bf}(r)| \ll |\xi^{Bf}(0)|$, the correction to the correlation function reduces to the simple approximate form

$$\Delta_{\xi}^{\xi^{AB}}(r) \sim \frac{\xi^{Af}(0)\xi^{Bf}(0)\bar{f}^2}{\sigma_f^2} = \bar{A}\bar{B}, \quad (2.7)$$

where $\bar{X} \equiv \langle WX \rangle / \langle W \rangle$ (with $X \in \{A, B\}$) is the mean value of X evaluated over the unmasked area. If the correlation is instead defined after subtracting the means over the unmasked areas, the limiting value is instead zero as expected.

For r sufficiently small that the two points are almost surely both unmasked or inside the same mask hole, with $f(\mathbf{x}') \approx f(\mathbf{x})$, we have the other limiting form

$$\Delta_{\xi}^{\xi^{AB}}(r) \sim \frac{\xi^{Af}(0)\xi^{Bf}(0)}{\sigma_f^4} (\bar{f}^2 - \sigma_f^2), \quad (2.8)$$

where $\bar{f}^2 \equiv \langle f^2W \rangle / \langle W \rangle$ is the mean value of f^2 evaluated over the unmasked area. If the mask systematically removes peaks of f , so that $\bar{f}^2 < \sigma_f^2$, and A and B have positive correlation to the foreground f , this is negative. If the means over the unmasked areas are subtracted before calculating the correlation functions, this becomes

$$\Delta_{\xi}^{\xi^{AB,0}}(r) \sim \frac{\xi^{Af}(0)\xi^{Bf}(0)}{\sigma_f^4} (\bar{\sigma}^2 - \sigma_f^2), \quad (2.9)$$

where $\bar{\sigma}^2 \equiv \overline{f^2} - \bar{f}^2$ is the point variance estimated over the unmasked area.

C. Poisson point sources

A Poisson model that describes the masking due to radio sources can also be handled analytically assuming a Gaussian distribution for the mean Poisson density. Following Paper I, we introduce the Poisson intensity field $\lambda(\mathbf{x})$, which defines the probability of observing no source in a given area \tilde{A} as $\exp(-\int_{\tilde{A}} d^2y \lambda(\mathbf{y}))$. On each location of a source, a disk of radius R is masked, and the mask W consists of the collection of these disks.

Under the assumption of a Gaussian λ , we may then write

$$\begin{aligned} \Delta \xi^{AB}(r) + \xi^{AB}(r) \\ = \frac{\langle A(\mathbf{x})B(\mathbf{x}') \exp(-\int_{\tilde{A}(\mathbf{y})} \lambda(\mathbf{y})d^2\mathbf{y}) \rangle}{\langle \exp(-\int_{\tilde{A}(\mathbf{y})} \lambda(\mathbf{y})d^2\mathbf{y}) \rangle}. \end{aligned} \quad (2.10)$$

In this equation, $\tilde{A}(\mathbf{y})$ is the area $D_R(\mathbf{y}-\mathbf{x}) + D_R(\mathbf{y}-\mathbf{x}') - D_R(\mathbf{y}-\mathbf{x})D_R(\mathbf{y}-\mathbf{x}')$ defined by the overlap of two disks D_R of radii R centred on \mathbf{x}' and \mathbf{y}' . Taking $A(\mathbf{x})$, $B(\mathbf{x}')$, and the integral over $\lambda(\mathbf{y})$ to be three correlated Gaussian variables, the Gaussian integral can be done analytically in terms of integrals over $\xi^{A\lambda}(r)$ and $\xi^{B\lambda}(r)$. These integrals can be evaluated directly numerically, or the result can be rewritten compactly in the form

$$\begin{aligned} \Delta \xi^{AB}(r) = [(\xi^{Af}(r) + \xi^{Af}(0) - ((\xi^{A\lambda} \cdot D_R) \star D_R)(r))] \\ \cdot [(\xi^{Bf}(r) + \xi^{Bf}(0) - ((\xi^{B\lambda} \cdot D_R) \star D_R)(r))], \end{aligned} \quad (2.11)$$

where $f(\mathbf{x})$ is the Poisson intensity $\lambda(\mathbf{x})$ convolved with D_R . Here, the $(\xi^{A/B\lambda} \cdot D_R) \star D_R$ terms are the convolution (\star) of the disk D_R with the disk-truncated correlation $D_R(r)\xi^{A/B\lambda}(r)$, which reduces to $\xi^{A/Bf}(0)$ for $r \ll R$ and vanishes for distances $r > 2R$. Each full term therefore changes smoothly from $\xi^{A/Bf}(0)$ at $r \sim 0$ to $\xi^{A/Bf}(0) + \xi^{A/Bf}(r)$ at $r \geq 2R$.

Note that $\xi^{Xf}(0) = -\bar{X}$, so if the means \bar{A} and \bar{B} over the unmasked areas are subtracted, this amounts to removing the $\xi^{A/Bf}(0)$ terms in Eq. (2.11). The entire correction is fourth order in the fields and hence small in either case. If there is a nonzero bispectrum, which is only third order in the fields, the actual real-world effect may be dominated by the non-Gaussian correlations rather than the very small Gaussian prediction.

III. SIMULATIONS

The analytic models described in the previous section were highly idealized. To study the impact of a correlated mask on real data, it is necessary to simulate CMB lensing

and correlated foreground fields in a realistic way. To do this, we followed Paper I and used the publicly available WEBSKY² simulation suite [59]. WEBSKY models the evolution of the matter distribution using the mass-Peak Patch method [60] in a volume of ~ 600 (Gpc/h)³ with $\sim 10^{12}$ particles over the redshift interval $0 < z < 4.6$. Despite being an approximate N-body method, mass-Peak Patch has been shown to reproduce the clustering properties of halos with good accuracy for both two-point and three-point statistics and their covariances across [61–63]. The simulation only includes dark matter particles, so we neglected baryonic effects in the analytical modeling discussed in the previous section as well. This realization of the matter distribution is used to produce full-sky maps of the CMB lensing convergence κ as well as extragalactic foregrounds based on an analytic halo model calibrated on existing observations and hydrodynamical simulations. The public release contains maps of tSZ and CIB at multiple frequencies, catalogs of radio sources based on the model of Refs. [64,65], and catalogs of the dark matter halos identified in the cosmological simulation.

A. Lensed CMB and CMB lensing

Following Paper I, we produced two sets of lensed CMB simulations that are later used to study the impact of correlated masking on lensing reconstruction. These sets share the same 100 unlensed CMB realizations, but they are lensed either using the same fixed deflection field constructed from the WEBSKY κ simulation (NG set) or using a different Gaussian random realization of the deflection field for each map (G set), where the deflection field has the same angular power spectrum as the realization of the WEBSKY κ map. We used the NG set to isolate the bias as it would appear on real data, while the G set was used to compute the error bars of our measurements and calibrate the lensing reconstruction normalization and $N^{(i)}$ biases. As such, the error bars displayed in the figures do not include any non-Gaussian contribution to the covariance. In the following, unless stated otherwise, error bars (or bands) displayed in figures are standard deviations of the plotted variable measured over the set of G simulations. An additional G-like set was computed to estimate the mean field of the quadratic estimator in order to debias the reconstructed lensing potential power spectrum.

The WEBSKY lensing κ map is constructed using the Born approximation and therefore neglects the effect of post-Born lensing, which is expected to significantly modify the overall non-Gaussianity in the CMB lensing κ [25] and thus the shape of the $N_L^{(3/2)}$ bias in the reconstructed lensing field [18,27]. Investigating the impact of post-Born effects on the mask bias would require us to include post-Born

²https://mocks.cita.utoronto.ca/index.php/WebSky_Extragalactic_CMB_Mocks.

lensing in the lensed CMB and (Born) lensing of the extragalactic foregrounds to avoid introducing spurious decorrelation of photon deflections along the line of sight (see, e.g., the discussion in Refs. [18,66]). As we do not have a similar set of simulations at our disposal, we investigated the role of post-Born effects only for masks built on the κ field using the Born and post-Born CMB lensing maps of Ref. [67].³ We used those maps to produce two additional sets of simulations where we lensed the same 100 realizations of unlensed CMB maps with deflection fields constructed from the Born and post-Born κ maps following the steps adopted for the WEBSKY κ map. We refer to these sets of simulations as BNG and the pBNG, respectively. To calibrate the normalization of the estimator, we also created a set of Gaussian simulations, as done for the G set, using the power spectrum of the post-Born κ map.⁴

B. Foreground masks

Using the WEBSKY simulations, we constructed three qualitatively distinct kinds of lensing-correlated masks:

- (i) W_{halo} masks, which roughly mimic masking of objects detected by a mass-limited cluster sample, for example, objects detected in a tSZ cluster survey with a given experimental noise level. We created different masks by selecting all the halos in the WEBSKY halo catalog with a mass above a certain mass threshold, M^{cut} defined in terms of $M_{500,c}$.⁵ For all the selected halos, we masked a disk centred on the halo position with a radius that is a multiple n of the $\theta_{500,c}$ halo angular size. In the following, we will adopt $n = 2$ as our default setup and will show results for $M^{\text{cut}} = (1.0, 1.8, 3.0) \cdot 10^{14} M_{\odot}/h$. For reference, Fig. 2 shows the distributions of angular size $\theta_{500,c}$ in the WEBSKY halo catalog.
- (ii) Foreground intensity threshold masks W_{κ} , W_{CIB} , and W_y . These masks are created by thresholding the WEBSKY κ , CIB at 217 GHz, and the tSZ Compton- y parameter maps, respectively, such that all the pixels above a specific value are removed from the analysis. We used different thresholds that effectively mask different sky fractions, $f_{\text{sky}}^{\text{mask}} = 0.6\%, 2.3\%, 6.7\%$. Before the thresholding step, we smoothed the foregrounds maps with a

³These adopt a *Planck* 2013 cosmology which is slightly different from the one used in the WEBSKY simulations.

⁴As noted in [27] differences due to post-Born correction in $C_L^{\kappa\kappa}$ are negligible for the purpose of this work.

⁵In the following, we usually define halo masses in terms of spherical overdensity masses in terms of $M_{500,c}$ and $M_{200,m}$. The first is the mass contained within the radius $R_{500,c}$ inside of which the mean interior density is 500 times the critical density. $M_{200,m}$, conversely, is the mass contained within the radius $R_{200,m}$ inside of which the mean interior density is 200 times the mean matter density of the universe.

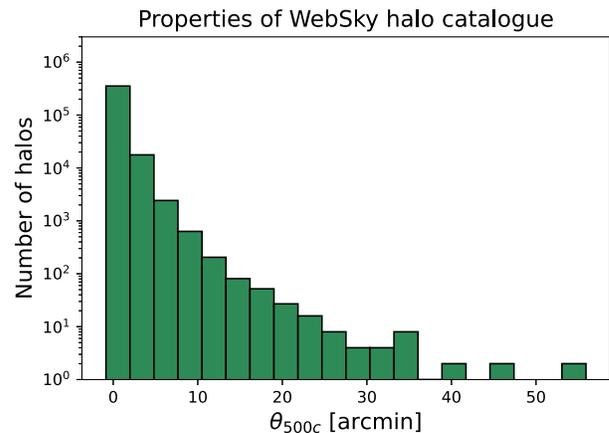


FIG. 2. Histogram showing the angular size, $\theta_{500,c}$, distribution of halos in the WEBSKY catalog for selected halos with $M_{500,c} > 10^{14} M_{\odot}/h$. This is consistent with the mass limit of detectable tSZ clusters in an S4-like survey (see Ref. [43] for more details).

Gaussian beams of full width at half-maximum ($\theta_{1/2}$) of 5.1' or 1.7'. This gives masks with more regular and connected holes, as expected from an experiment with a comparable beam size. By construction, these masks allow us to explore the biases for foregrounds with different degrees of correlation with CMB lensing, where W_{κ} is a limiting-case where the masked field is 100% correlated, while the CIB map is correlated at the $\gtrsim 70\%$ level at $\ell < 1000$ and the tSZ map is correlated at the 30%–50% level. The W_{halo} mask is strongly correlated to the W_y mask, since the main peaks of the y -parameter map are associated with massive clusters.

- (iii) W_{rs} masks remove resolved radio point sources (and are the same as those described in Paper I). For this purpose, we selected all the sources with a measured flux above the 5σ detection limit for *Planck*, SO, and S4 and cut out a circular region of the sky around the source. The radius of these holes has been chosen to be $2\theta_{1/2}$ of each frequency channel. Each of these masks is built as the product of three masks obtained by selecting sources in the three frequency bands ranging from 90 to 225 GHz most relevant for small-scales CMB power spectra measurements (for more details, see Table I of Paper I).

To isolate the impact of mask correlations, we apply to each mask a single random rotation to build a new mask, W^{rot} , which is effectively uncorrelated with CMB lensing but retains all the other nontrivial mode-coupling effects due to cut sky and hole shapes (we neglect the small area of residual correlation around the poles of the random rotation axis). In Fig. 3, we show a cutout of the full sky κ and tSZ y -parameter maps from the WEBSKY simulation together with an example of the masks used in our analysis.

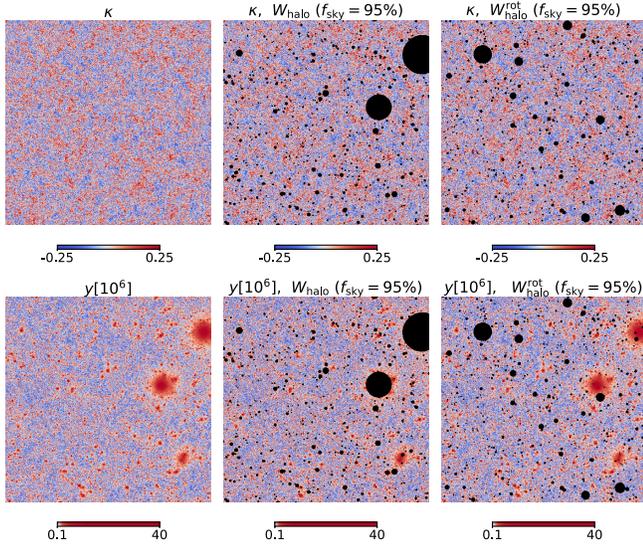


FIG. 3. $20 \times 20 \text{ deg}^2$ cutout of the κ (top) and tSZ y parameter maps (bottom) of the WEBSKY simulation suite. We highlight the masked pixels in black for one of the W_{halo} masks employed in this work. The left column displays the unmasked κ and y fields. The middle column shows the same fields masked with a mask that removes halos with $M > M^{\text{cut}} = 1.8 \times 10^{14} M_{\odot}/h$. As is evident from the bottom row plot, the masked areas typically correspond to dense areas of strong tSZ emission. The right column shows κ and y masked with a mask uncorrelated with the underlying lensing field constructed randomizing the position of the halos used for the W_{halo} mask.

C. Biases from direct κ masking

To measure the effect on simulations, we applied the foreground masks described in Sec. III B, W_X , to the WEBSKY κ map, and estimated the power spectrum over the unmasked area by deconvolving the effect of the mask using the MASTER [68] algorithm as implemented in the publicly available NAMASTER⁶ package [69]. We then compared this mask-deconvolved power spectrum with the angular power spectrum of the WEBSKY κ computed on the full sky. The results are shown in Fig. 4. For each spectrum, we adopted a variable binning in L : $\Delta_L = 1$ until $L_{\text{max}} = 10$, 4 bins of $\Delta_L = 20$ ($L_{\text{max}} < 100$), 4 bins of $\Delta_L = 100$ ($L_{\text{max}} < 500$), 14 bins of $\Delta_L = 250$ ($L_{\text{max}} < 4000$), and $\Delta_L = 1000$ for $L_{\text{max}} > 4000$. We estimated the error bars on the measured bias using a jackknife approach. For this purpose, we divided the sky map into $N_{\text{patches}} = 48$ subregions of equal areas corresponding to pixels of a HEALPIX pixelization with $n_{\text{side}} = 2$. For the i th subregion $W_{\text{cut}}^{(i)}$ and for each LSS-correlated mask W_X , we estimated the fractional difference between the $C_L^{\kappa\kappa}$ computed on the sky masked with $W_X \cdot W_{\text{cut}}^{(i)}$, and its value computed on the sky masked with

$W_{\text{cut}}^{(i)}$ only. The error bar is then given by the covariance of the fractional differences averaged over all the subregions rescaled by $N_{\text{patches}} - 1$ (see, e.g., Appendix B of Ref. [70]).

Since all our masks select areas of the sky where mass overdensities are present, the recovered power spectrum has less power compared to its full sky value. We do not observe this loss of power when we apply the W_X^{rot} masks that are uncorrelated with the lensing signal.

The right panel of Fig. 1 shows that the halo model describes reasonably well the effect observed when masking the WEBSKY κ map with W_{halo} as also shown in the top-right panel of Fig. 4; in particular, the analytical curves match well the shape of the biases observed in the simulations. Changing the halo mass function only gives fairly minor changes to the analytic predictions. We note that, since the halo mass functions and related bias model are usually formulated (or calibrated on simulations) in terms of $M_{200,m}$, while we perform cluster masking of κ in terms of $M_{500,c}$, we can compare the curves in the two plots noting that $M_{500,c} \approx 2M_{200,m}$.⁷ For the W_{halo} mask, reducing M^{cut} increases the number of objects that are masked and thus also the fraction of sky area that is masked. This leads to a progressively larger power deficit. Such suppression is more relevant at very large scales and at $L \sim 2000$ in broad agreement with the analytical predictions of Sec. II A. The shapes of the biases induced by W_{halo} and W_y are similar. This is due to the fact that, even though they are built in different ways, they are highly correlated as they mask similar objects and locations as discussed in Sec. III.

Proceeding clockwise in Fig. 4, we show the results for W_{rs} masks. All the relative differences between the $\kappa\kappa$ -spectrum computed on the masked sky and the $\kappa\kappa$ -spectrum computed on the full sky are quite small but still not compatible with zero, considering the error bars. The theoretical curves, obtained using Eq. (2.11), are not in agreement with the simulation results. This means that the Gaussian foreground peaks model is not completely able to reproduce the behavior of these radio source masks, consistent with non-Gaussian effects dominating the purely Gaussian predictions.

For the threshold masks, the second and third rows of Fig. 4, the dashed lines are the theory curves obtained using the pure Gaussian model described in Sec. II B, while the solid lines represent the semiempirical form of Eq. (2.6), where the spectra and cross-spectra of $W(\mathbf{x})$, $(fW)(\mathbf{x})$ and $(f^2W)(\mathbf{x})$ have been computed directly from our set of foreground masks and the WEBSKY κ , y , and CIB fields. Overall, the semiempirical Gaussian model using the actual masks and foreground fields seems to match better the data points, in particular for W_{κ} , even though, particularly for W_{CIB} and, even more, W_y masks, there is a clear

⁶<https://github.com/LSSTDESC/NaMaster>.

⁷We used the COLOSSUS code (<https://bitbucket.org/bdiemer/colossus/src/master/>) to perform an accurate conversion between the two mass definitions at a specific redshift in the integration.

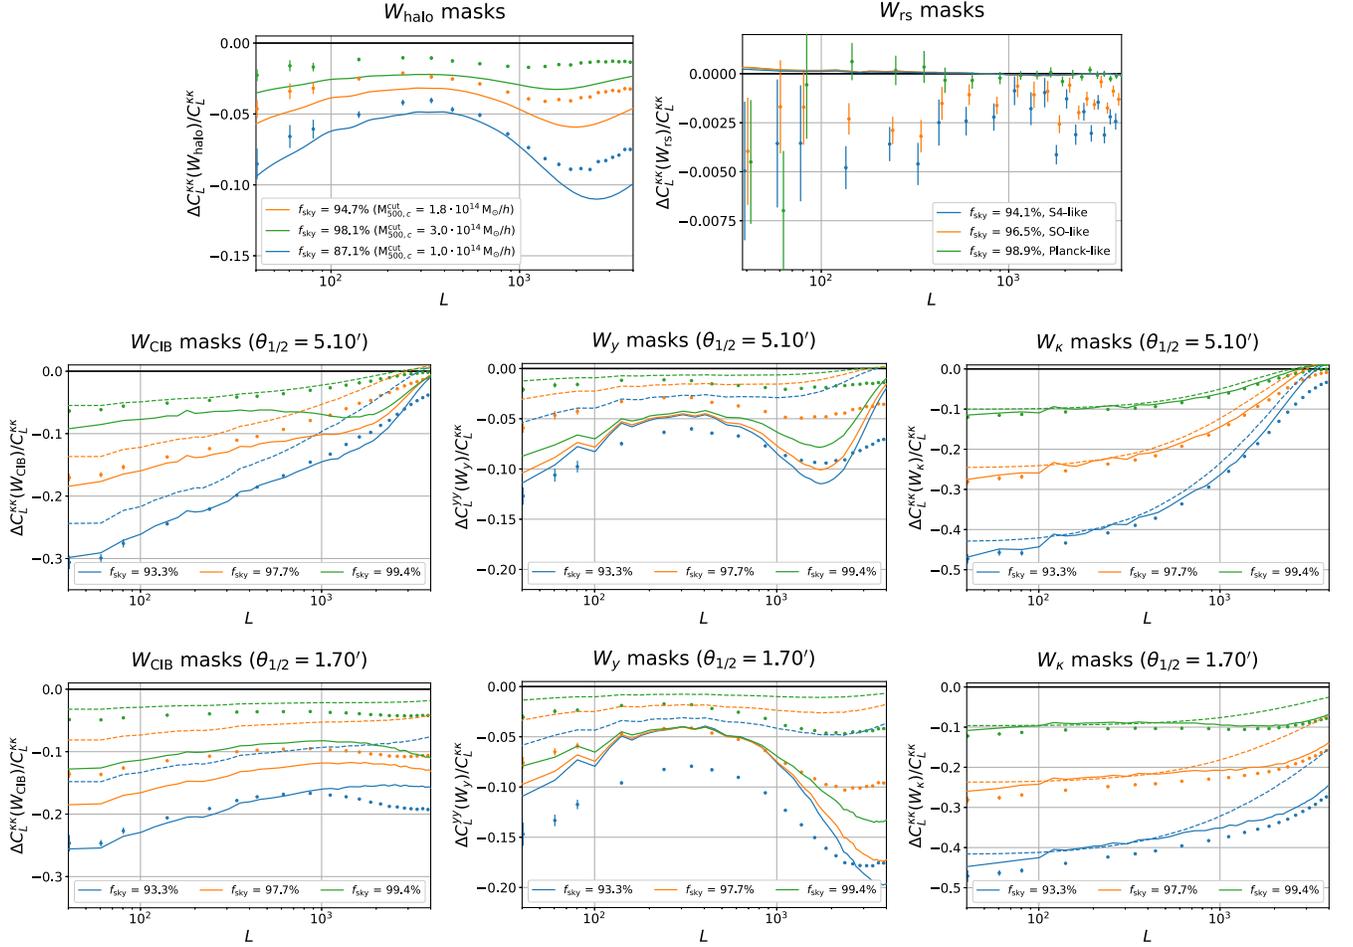


FIG. 4. Effect of LSS-correlated masks on the CMB lensing convergence power spectrum C_L^{KK} as a fractional difference between the C_L^{KK} computed on the masked sky and its value computed on the full sky. Each panel shows results for different masks. Different colors show the result for masks retaining different sky fractions after masking. The W_{halo} plot legend shows the mass limits of the halo samples used for the mask, and the W_{rs} plot legend shows the source count detection limit used to construct the radio-source mask for the various listed experiments. Simulation measurements are shown as data points. For W_{halo} and W_{rs} masks, the solid lines show the analytic predictions of Sec. II A and of Eq. (2.11), respectively. For the W_K , W_y , and W_{CIB} threshold masks, the semianalytic theory predictions are shown in solid, and the pure Gaussian model is shown as dashed. For the former, we evaluated the expectation values in Eq. (2.6) empirically from our simulations, while the latter uses a fully analytic model. The middle and bottom rows show results obtained for different smoothings of the foreground fields.

disagreement at small scales. This is not surprising as κ is the most Gaussian of the three foreground fields considered here. The CIB contains a significant shot-noise contribution from individual bright infrared galaxies and the y from highly collapsed galaxy clusters that induce significant non-Gaussianities in the maps. We investigate this disagreement for the threshold masks in Appendix B.

IV. LENSING RECONSTRUCTION ON MASKED FIELDS

In Sec. III C, we have directly masked the CMB lensing convergence to roughly quantify the bias induced by the foreground masks described in Sec. III B. However, the true κ field is not a direct observable and has to be reconstructed

from the observed CMB maps potentially also in combination with other external matter tracers. Hence, in this section, we quantify the impact of LSS-correlated masks, removing the corresponding regions in lensed CMB maps that are then used to perform lensing reconstruction using a quadratic estimator. For this first analysis, we only use the CMB temperature field, since the foreground contamination is less important for polarization and the properties of the polarized sources are currently much less well understood [71].

For the reconstruction of the lensing potential, we mainly followed the same steps as the *Planck* lensing pipeline described in Ref. [72]. For CMB temperature lensing reconstruction, this is just an optimized and generalized version of the lensing quadratic estimator (QE) of Ref. [73].

We used the reconstruction pipeline implemented in the public PLANCKLENS code,⁸ and we refer the reader to Ref. [72] for more details. The procedure can be summarized in four steps: (1) optimal filtering of input “data” CMB maps; (2) construction of the lensing quadratic estimator; (3) mean-field subtraction and normalization of the lensing estimate; and (4) computation of the lensing power spectrum, subtraction of its additive biases, and Monte Carlo correction of its normalization.

To measure the mask biases in lensing reconstruction, we used the two sets of Monte Carlo simulations of lensed CMB realizations describe in Sec. III A. We used the NG set to isolate the bias as it would appear on real data, while the G set was used to compute the mean-field, the realization-dependent $N^{(0)}$ (RD- $N^{(0)}$) noise bias, and the multiplicative Monte Carlo correction assuming no mask-lensing correlation. We also used the G set to estimate error bars on the auto- and cross-spectra estimators described in the following. Since the WEBSKY suite includes only a single realization of κ , our numerical results based on the NG set are limited on large scales by the cosmic variance of this fixed realization of the lensing field.

In the following, we considered two experimental setups representative of an SO-like and an S4-like survey. For SO, we assumed an effective white noise level in temperature of $\sim 6.7 \mu\text{K-arc min}$ and $\theta_{1/2} = 1.5'$, consistent with the publicly available effective baseline noise configuration after component separation.⁹ For the S4-like survey, we assumed a beam with $\theta_{1/2} = 1.0'$ and isotropic uncorrelated $\sim 1 \mu\text{K-arc min}$ noise for temperature.¹⁰ The results shown in the following are computed using SO-like experimental specifications. In addition, we considered an S4-like experimental setup for a subset of foreground masks: the W_{rs} mask with $f_{\text{sky}} = 94.1\%$, which includes all the sources with a measured flux above the detection limit for S4; the W_{halo} with $M^{\text{cut}} = 10^{14} M_{\odot}/h$, which removes a mass-limited tSZ-selected cluster sample for an S4-like survey [74]; and the extreme cases of W_{CIB} and W_y with $f_{\text{sky}} = 93.3\%$ and $\theta_{1/2} = 1.7'$.

After adding a realization of the isotropic noise, the input maps are then masked using all the unapodized W_X masks discussed in Sec. III. The (optimal) filtering step produces Wiener-filtered maps that provide the minimum variance estimate of the full-sky lensed CMB based on the information in the unmasked area. Optimal filtering [72,75,76] is particularly valuable with this kind of mask compared to more basic (but faster) inverse-variance-weighted or

isotropic filtering. Figure 5 shows how the optimal filtering operation is able to fill back some information inside small masked regions, effectively recovering information that was masked (assuming no residual foregrounds outside the masked area). This increases the information available, reduces any complications due to sharp mask cuts, and because the masked area is effectively reduced can substantially reduce biases when the mask is correlated to the signal.

To build the filtered CMB maps, we used a fiducial lensed CMB temperature spectrum including multipoles $100 \leq \ell \leq 4000$. Removing multipoles $\ell < 100$ generates little loss of information for lensing reconstruction. Since we approximate the noise as white and isotropic, the noise in the filter is also taken to be isotropic and consistent with the simulated value.

The real-space lensing deflection estimator is then built from a pair of filtered maps discussed above. The gradient part¹¹ of the quadratic lensing deflection estimator, \hat{g}_{LM} , contains information on the lensing potential, which is estimated using [21]

$$\hat{\phi}_{LM} \equiv \frac{1}{\mathcal{R}_L^{\phi}} (\hat{g}_{LM} - \langle \hat{g}_{LM}^{\text{MF}} \rangle), \quad (4.1)$$

where \mathcal{R}_L^{ϕ} is the nonperturbative response function defined to make the lensing reconstruction unbiased on the full sky [18,77,78] and $\langle \hat{g}_{LM}^{\text{MF}} \rangle$ is the mean field of the estimator. Note that we are constructing our lensing estimators in the hypothesis of no lensing-mask correlation, as appropriate for quantifying the bias on standard methods. We compute the mean field using the simulations of the G set rather than, for example, defining the mean field by averaging over simulations with lensing fields correlated to the fixed mask. Since the mean field depends on the mask, for lensing-correlated masks the mean field is actually correlated to the true lensing potential ϕ .

In our analysis, we make two estimates of the (uncorrelated) mean field for each mask using two different independent sets of independent 25 G simulations, which we denote $\langle \hat{g}_{LM}^{\text{MF}} \rangle_1$ and $\langle \hat{g}_{LM}^{\text{MF}} \rangle_2$, respectively (using 50 G simulations in total, 25 for each mean field). The lensing power spectrum is estimated by cross-correlating two lensing map estimates $\hat{\phi}_1$ and $\hat{\phi}_2$ following Eq. (4.1), where we subtracted $\langle \hat{g}_{LM}^{\text{MF}} \rangle_1$ and $\langle \hat{g}_{LM}^{\text{MF}} \rangle_2$, respectively, to avoid reconstruction noise in the mean-field cross-correlation. The estimate of the lensing power spectrum of the single realization is then obtained as

¹¹The curl component is expected to be zero to a good approximation and is zero in the WEBSKY lensing field by construction.

⁸<https://github.com/carronj/plancklens>.

⁹Details of the noise model for SO can be found at https://github.com/simonsobsso_noise_models. Note that the beam size is not entirely consistent with the $1.7'$ fiducial foreground smoothing that we use, which was chosen to match the smoothing used in Paper I.

¹⁰Details can be found at https://cmb-s4.uchicago.edu/wiki/index.php/Survey_Performance_Expectations.

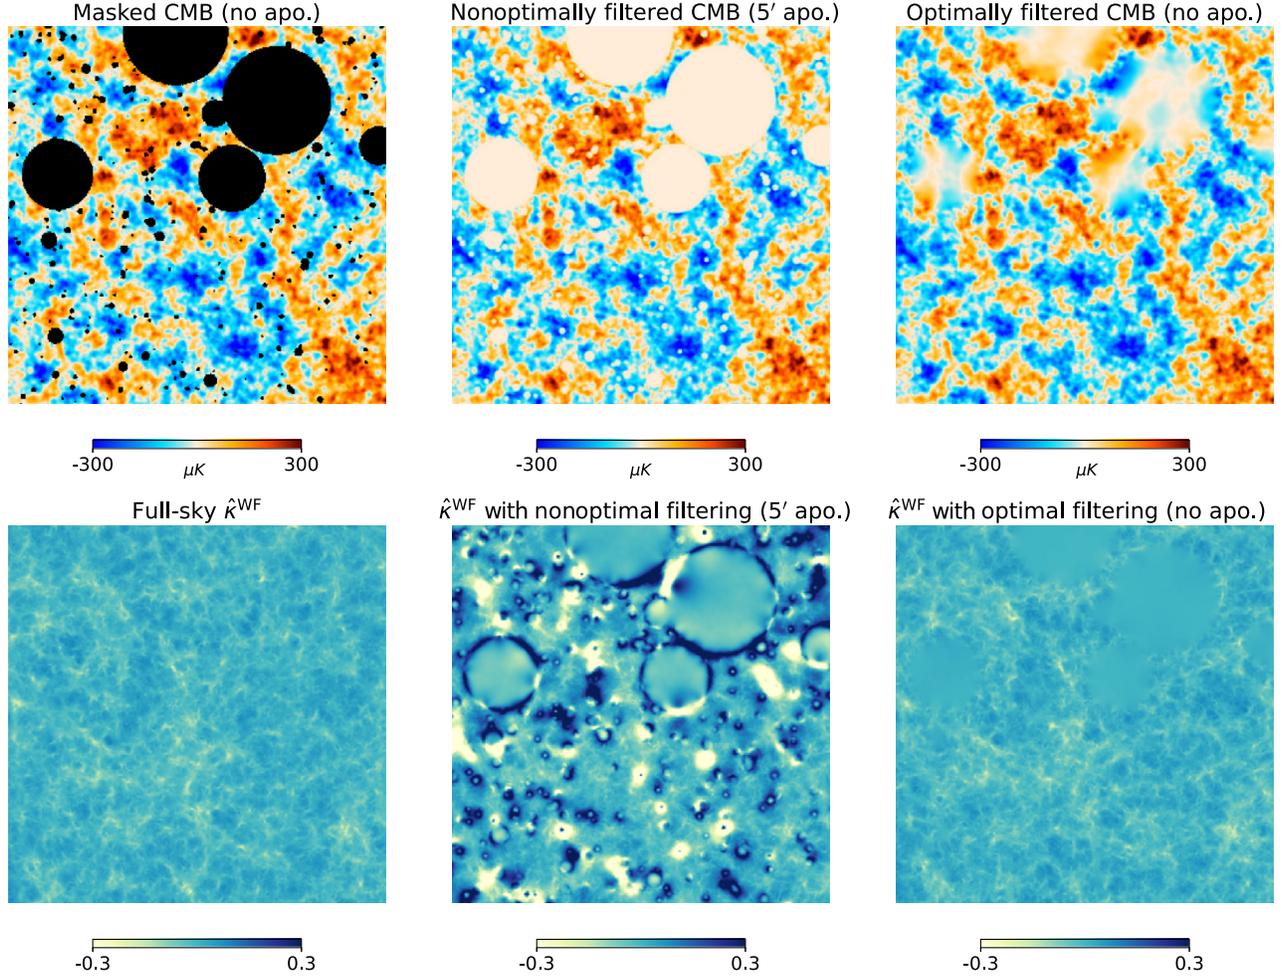


FIG. 5. Comparison between optimal and nonoptimal Wiener filtering. Top panel, from left to right, shows the masked CMB field (W_{halo} with $f_{\text{sky}} = 94.7\%$), the filtered CMB field using a nonoptimal isotropic method, and the filtered CMB field using optimal filtering as described in Refs. [72,75]. Bottom panel, from left to right, shows the Wiener-filtered full-sky reconstruction κ map and the filtered reconstructed κ map using nonoptimal and optimal filtering. In the nonoptimal filtering case, the mask has been apodized using the C^2 function (effectively a cosine) implemented in NAMASTER with an apodization scale of $5'$. This helps to slightly reduce the ringing effects, but large hole-induced reconstruction noise edge effects still dominate. In all the κ maps, a significant fraction of the small-scale structure is the non-Gaussian reconstruction noise.

$$\hat{C}_L^{\hat{\phi}_1 \hat{\phi}_2} \equiv \frac{1}{(2L+1)f_{\text{sky}}} \sum_{M=-L}^L \hat{\phi}_{1,LM}^* \hat{\phi}_{2,LM}, \quad (4.2)$$

where $f_{\text{sky}} = \sum_p W_{Xp}/N_{\text{pix}}$ is the unmasked sky fraction.

From the above estimator, we then subtract the realization-dependent estimate of the Gaussian (disconnected) lensing bias, $\text{RD-}N_L^{(0)}$. Subtracting this term, together with the mean-field subtraction, has the effect of removing the disconnected signal expected from Gaussian fluctuations even in the absence of lensing [77,79]. The $\text{RD-}N_L^{(0)}$ bias is defined as

$$\text{RD} - N_L^{(0),d} \equiv \langle 4\hat{C}_L^{di} - 2\hat{C}_L^{ij} \rangle, \quad (4.3)$$

where angle brackets denote an average over pairs of distinct i, j G simulations. Here, \hat{C}_L^{di} is the estimator in Eq. (4.2) with $\hat{\phi}_1$ reconstructed from the d th CMB realization of the NG set and $\hat{\phi}_2$ using the i th simulation of the G set; \hat{C}_L^{ij} is instead the estimator of Eq. (4.2) with $\hat{\phi}_1$ and $\hat{\phi}_2$ reconstructed from the i th and the j th simulations of the G set, respectively. For each mask (both LSS-correlated and uncorrelated), $\text{RD-}N^{(0),d}$ is calculated for each NG “data” simulation, d , using 20 different pairs of independent G simulations.

Figure 6 shows a typical mean-field power spectrum corrected for the nonperturbative response and an example of the $\text{RD-}N^{(0),d}$.

For a given lensing reconstruction field, we can also define the cross-spectrum estimator with the true field

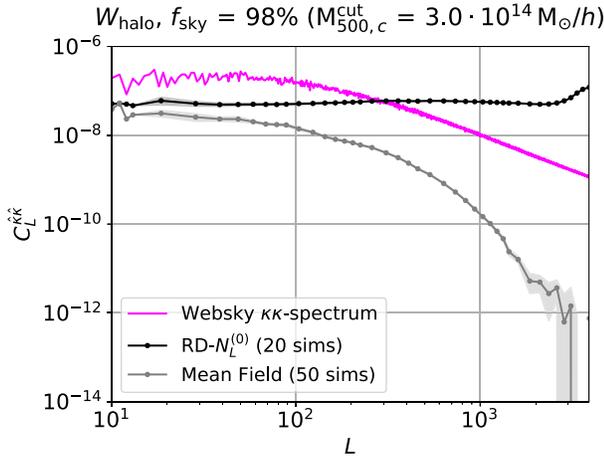


FIG. 6. The $RD-N_L^{(0)}$ bias [black, as described in Eq. (4.3)] and mean-field power spectrum corrected for the response (gray) for the case of the fixed halo mask described in the plot title. We show (in purple) the WEBSKY full-sky $C_L^{\kappa\kappa}$ for comparison.

$$C_L^{\hat{\phi}\phi} \equiv \frac{f_L^{\text{MC}}}{(2L+1)f_{\text{sky}}} \sum_{M=-L}^L \hat{\phi}_{LM}^* \phi_{LM}, \quad (4.4)$$

where $\hat{\phi}$ is the lensing potential estimate of Eq. (4.1), ϕ is the true lensing potential field, and f_L^{MC} is a Monte Carlo (MC) correction defined to make the estimator unbiased for Gaussian lensing potential uncorrelated with the mask. Specifically, we define

$$\frac{1}{f_L^{\text{MC}}} \equiv \left\langle \frac{[C_L^{\hat{\phi}\phi}]^{-1}}{(2L+1)f_{\text{sky}}} \sum_{M=-L}^L \hat{\phi}_{LM}^* \phi_{LM} \right\rangle_{50 \text{ G sims}}. \quad (4.5)$$

In practice, this ratio is calculated for the binned spectra rather than individual L , using the binning scheme described in Sec. III C. Note that for large (e.g., galactic) masks f_{sky} is a reasonable approximate normalization so that $f_L^{\text{MC}} \approx 1$. However, when the mask contains a large number of small holes due to, e.g., point source masking, the f_{sky} correction becomes a much worse approximation as effectively much less mask area is lost after optimal filtering and reconstruction. We discuss this aspect in more detail in Appendices A and A 3. Here, the inclusion of f_{sky} does not affect the result, since it simply amounts to a redefinition of f_L^{MC} .

We also define the lensing power spectrum estimator

$$C_L^{\hat{\phi}_1 \hat{\phi}_2, \text{RD}} = (f_L^{\text{MC}})^2 (\hat{C}_L^{\hat{\phi}_1 \hat{\phi}_2} - \text{RD} - N_L^{(0)}), \quad (4.6)$$

where f_L^{MC} is defined as above. However, the $(f_L^{\text{MC}})^2$ normalization calibrated on the cross-spectrum may not be the correct normalization factor for the autospectrum. Moreover, even in presence of Gaussian lensing fields and

LSS-uncorrelated masks, the estimator in Eq. (4.6) does not provide an unbiased estimate of the CMB lensing power spectrum as it still retains the $N_L^{(1)}$ noise bias induced by signal-dependent contractions [23]. For non-Gaussian lensing field, the estimator also retains the noise term involving the three-point function of the lensing field ($N_L^{(3/2)}$). Here, we regard $N_L^{(1)}$ and $N_L^{(3/2)}$ as part of the signal contained in $C_L^{\hat{\phi}_1 \hat{\phi}_2, \text{RD}}$ and only consider differences between the estimated power spectrum (that includes these additional noise bias terms) when using LSS-correlated and uncorrelated masks.

We run the entire end-to-end estimation pipeline 20 times for each mask (both correlated and uncorrelated), taking the “data” each time to be one of the NG lensed CMB simulations. Results are then averaged over these simulations to reduce Monte Carlo noise from variations of the unlensed CMB. However, the cosmic variance of the lensing field is not reduced since all realizations in the NG set share the same single WEBSKY κ simulation that is currently available. To make a comparison with the results obtained in Sec. II, we plot results for the reconstructed convergence field $\hat{\kappa}$ instead of $\hat{\phi}$.¹²

The effect of the mean-field and $RD-N_L^{(0)}$ bias on the reconstructed $\hat{\kappa}$ -spectra is shown in the first row of Fig. 7 as relative differences between the reconstructed $\hat{\kappa}$ -spectrum and the true WEBSKY $\kappa\kappa$ -spectrum for three different W_{halo} masks. From left to right, the reconstructed auto-spectra are plotted without mean-field and $RD-N_L^{(0)}$ corrections, then subtracting the mean field only, and finally subtracting both mean field and the $RD-N_L^{(0)}$. The main correction to the reconstructed $\hat{\kappa}$ -spectra comes from subtracting the $RD-N_L^{(0)}$ bias. The rise on small scales is related to the unsubtracted $N_L^{(1)}$. The second row of Fig. 7 shows the effect of the mean-field subtraction on the relative differences between the cross spectrum $\hat{\kappa}\kappa$ and the true WEBSKY $\kappa\kappa$ -spectrum. Note that for correlated masks, the mean field is also important for the cross-correlation since the mask, and hence the mean field, is correlated to the signal. The bias induced by the mask on the $\hat{\kappa}\kappa$ spectra is strongly reduced when we remove the mean-field term, especially at intermediate scales where the difference is then dominated¹³ by $N_L^{(3/2)}$. In all the plots of Fig. 7, the full-sky cases are shown in purple as reference. Similar results are obtained for all the other masks described in Sec. III B.

¹²Note that $\hat{\kappa}_{LM} \equiv \frac{1}{2}L(L+1)\hat{\phi}_{LM}$.

¹³Note that the $N_L^{(3/2)}$ signal is overestimated because the WEBSKY simulations do not include post-Born lensing, which largely has an opposite sign.

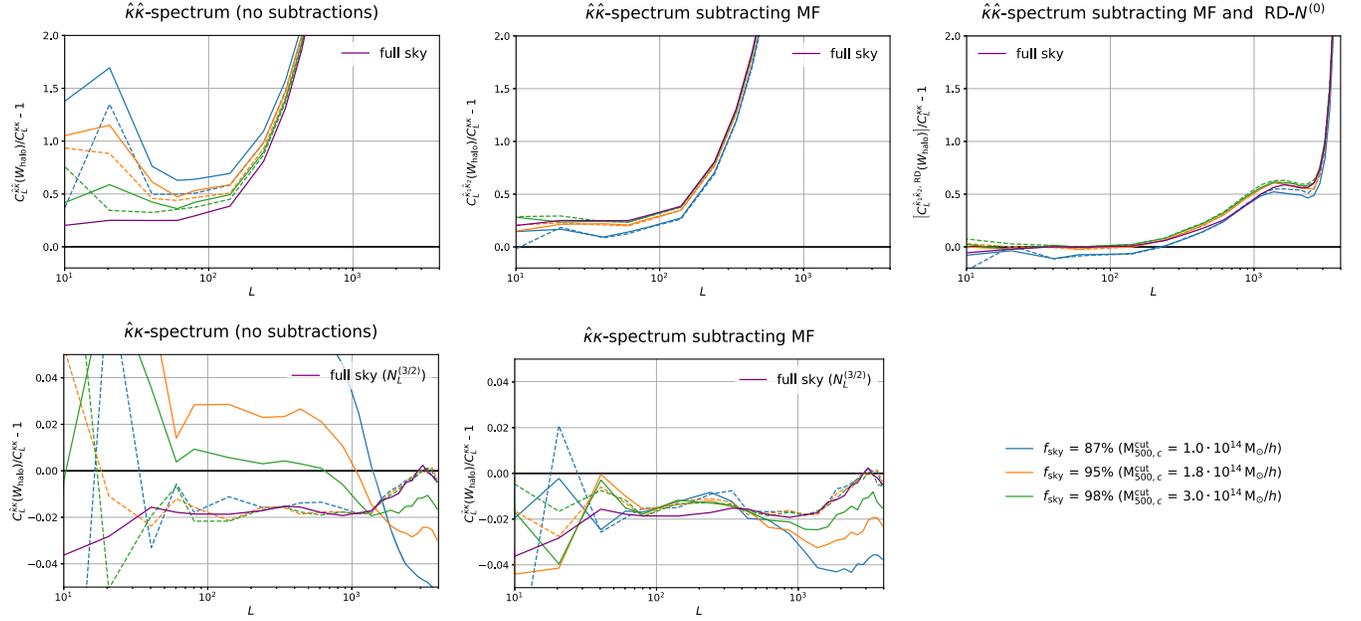


FIG. 7. Top: the left panel shows the raw reconstructed lensing autospectra, while the middle and right panels show the reconstructed autospectra after subtracting the mean field alone and the mean-field and the RD- $N_L^{(0)}$ noise bias respectively. Bottom: the reconstructed cross-spectra with and without the mean-field subtraction are shown on the left and right sides, respectively. All these curves include the Monte Carlo normalization correction, f_L^{MC} . The results are obtained using the W_{halo} masks for different values of the mass detection threshold, M^{cut} , whose corresponding f_{sky} is reported in the legend. The purple lines correspond to the full-sky analysis, where, for the cross-correlation, the final difference is consistent with $N_L^{(3/2)}$ reconstruction bias and for the autospectrum is dominated by $N_L^{(1)}$. Dashed-lines show the results obtained using the LSS-uncorrelated masks, $W_{\text{halo}}^{\text{rot}}$.

V. NUMERICAL RESULTS

A. CMB lensing power spectrum

To isolate the effects due to the correlation between the mask and the lensing field, we computed the difference between the two-point correlation function of the reconstructed CMB lensing obtained with the W_X masks and the one obtained with the rotated uncorrelated masks, W_X^{rot} . The rotated results have no correlated mask effects but retain all the other nontrivial mode-coupling effects due to cut sky and hole shapes, as well as $N_L^{(3/2)}$ and $N_L^{(1)}$ to the extent that they are not modified by an LSS-correlated mask.

We show the results of these measurements in Figs. 8, 9, and 11 for the reconstructed auto- and cross-spectra respectively, showing the bias for all the masks, W_{halo} , W_κ , W_{CIB} , W_y , and W_{rs} . We estimated the error bars of these measurements computing the estimators of Eqs. (4.2) and (4.4) for both W_X and W_X^{rot} using as “data” independent sets of 20 G simulations. The errors are then taken as the standard deviations of the differences between correlated and uncorrelated (randomly rotated) mask results.

As expected, the biases become larger as we increase the masked fraction of the sky. However, the amplitude of the lensing reconstruction biases are significantly reduced with respect to those obtained from directly masking the lensing field presented in Sec. III C (see, e.g., Figs. 8 and 4 for a

direct comparison). The optimal filtering used by the lensing reconstruction pipeline substantially reduces the fraction of the lensing information that is removed by the mask, both because the filtering recovers some of the CMB modes inside the mask holes and because the lensing reconstruction itself is able to recover much of the information about lensing modes on scales larger than the hole size (see Appendix A for an analytic discussion).

The remaining biases induced by W_{halo} and W_y (with $\theta_{1/2} = 5.1'$) masks are mainly relevant on small scales. The S4-noise case considered for the W_{halo} mask shows a similar trend, with a remaining power spectrum bias of the 2%–5% level for $L \gtrsim 2000$. For the W_κ and W_{CIB} masks, the bias is roughly constant across all the scales and has a magnitude of $\sim 1\% - 10\%$ of the signal depending of the f_{sky} . The bias induced by W_κ is larger since this is the limiting case where the mask is 100% correlated with the κ field. The bias induced by W_{rs} is negligible for both SO and S4.

In Figs. 8 and 9, the foreground fields used to build the W_y , W_{CIB} , and W_κ masks were smoothed with a Gaussian beam of $\theta_{1/2} = 5.1'$ prior to thresholding, similar to the beam of current *Planck* data. This operation leads to masks with large connected holes on the scale of the smoothing. However, future experiments such as SO and S4 will observe the sky at higher angular resolution. To investigate

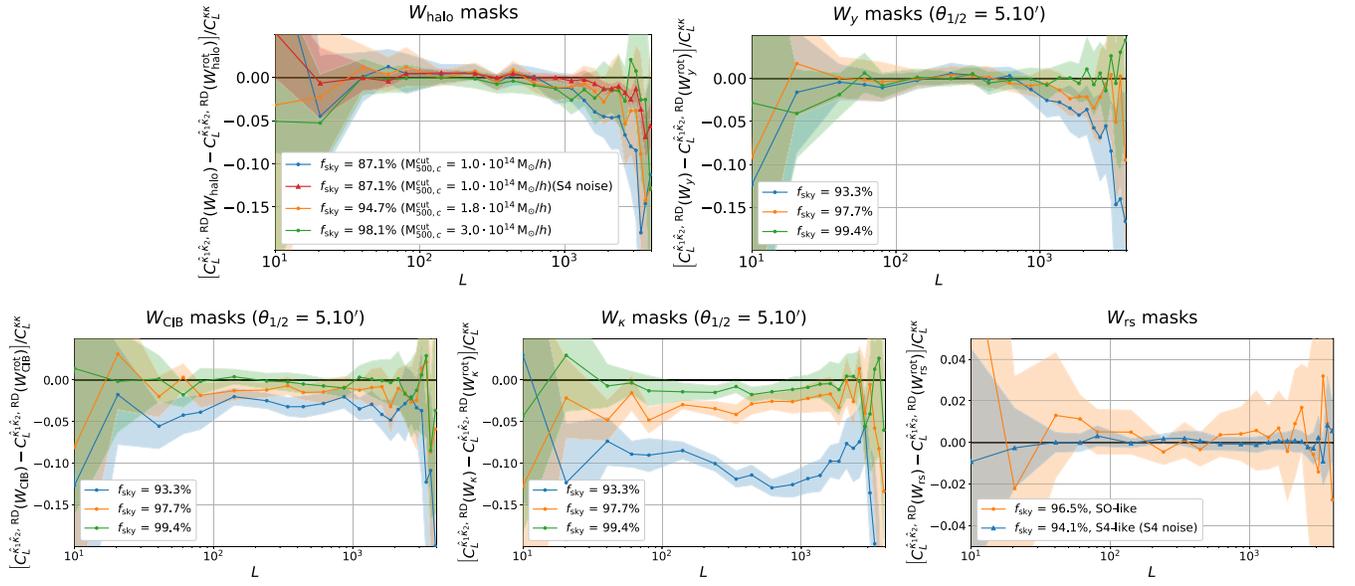


FIG. 8. Effect of LSS-correlated masking on the reconstructed CMB lensing convergence autospectrum obtained with the estimator of Eq. (4.6). The mask biases are computed as the differences between reconstructions performed with correlated and uncorrelated masks. The plots show the amplitude of the bias relative to the true lensing power spectrum. Different LSS-correlated masks are shown in different panels. The foreground fields have been smoothed with a $\theta_{1/2} = 5.1'$ Gaussian beam prior to thresholding. Unless stated otherwise, we assumed an SO-like CMB temperature noise in the lensing reconstruction.

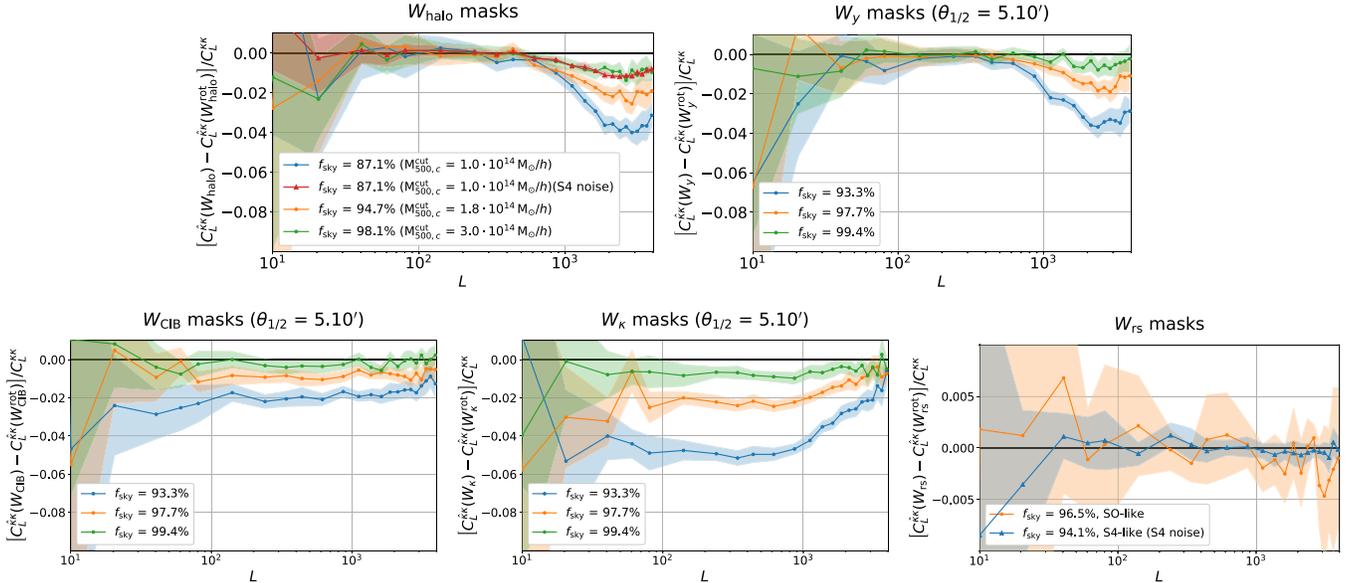


FIG. 9. Same as Fig. 8 for the cross-spectrum between the reconstructed CMB lensing convergence field and the true lensing field.

the sensitivity of our result to this smoothing scale, we performed a similar analysis on the masks obtained by smoothing the foreground field with a $\theta_{1/2} = 1.7'$ beam prior to thresholding. In Fig. 10, we show the comparison of the masks obtained with the two different smoothing scales. Figure 11 shows the measurements of the mask bias for this set of masks. From the direct-masking results obtained with these masks in Sec. III C, we would expect

larger biases at the level of the autospectra, especially at smaller scales. However, comparing Fig. 11 with the results shown in Figs. 8 and 9 using the larger smoothing scale, we actually see substantially smaller biases when we mask the same fraction of the sky. Reconstruction with optimal filtering is able to recover more information in the holes when we increase the number of holes but also to reduce their size. There is now only a small residual bias at

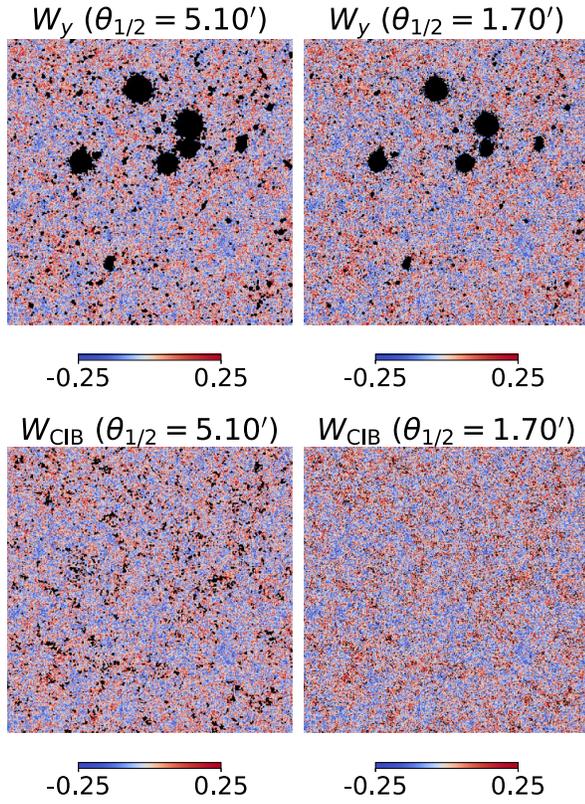


FIG. 10. Comparison between the masked CMB lensing convergence using W_y and W_{CIB} masks constructed from foreground fields smoothed with different values of $\theta_{1/2}$. While the shape of the masked regions changes significantly, the total masked sky fraction is the same ($f_{\text{sky}}^{\text{mask}} = 6.7\%$).

$L \gtrsim 1000$ for the W_y mask, where the signal is anyway largely dominated by reconstruction noise. The biases in the S4-noise case are instead completely negligible both for W_y and W_{CIB} . The reconstructed CMB lensing field can also be used to construct templates of the lensed CMB in order to perform delensing of the observed CMB data, in particular the B-mode polarization (see Refs. [80,81] for recent applications on data). As discussed in Ref. [82], an unbiased measurement of the lensing potential at $L \lesssim 1500$ would be sufficient to resolve with subpercent accuracy the lensing B-mode signal, which is the part of the signal for which the coupling induced by lensing is the most non-local. The mask biases on the reconstructed CMB lensing map and power spectra at these angular scales are small for the most realistic masks we considered, and even more for the S4-like observations we considered, with high-resolution and low-noise, a regime for where delensing would bring larger improvements for cosmological constraints. As such, we do not expect the mask biases to become a significant problem for CMB internal delensing for future datasets.

To estimate the impact of post-Born lensing on the correlated mask biases, we ran the entire end-to-end mask bias estimation pipeline, as described in Sec. IV, using the pBNG and BNG simulations, computing f_L^{MC} corrections and the mean field of the QE on the pBG simulation set. For both the pBNG and BNG simulations, we built a new correlated foreground mask, thresholding the corresponding κ field such that $f_{\text{sky}}^{\text{mask}} = 6.7\%$. This limiting case of a κ mask is the only case we considered. Post-Born lensing

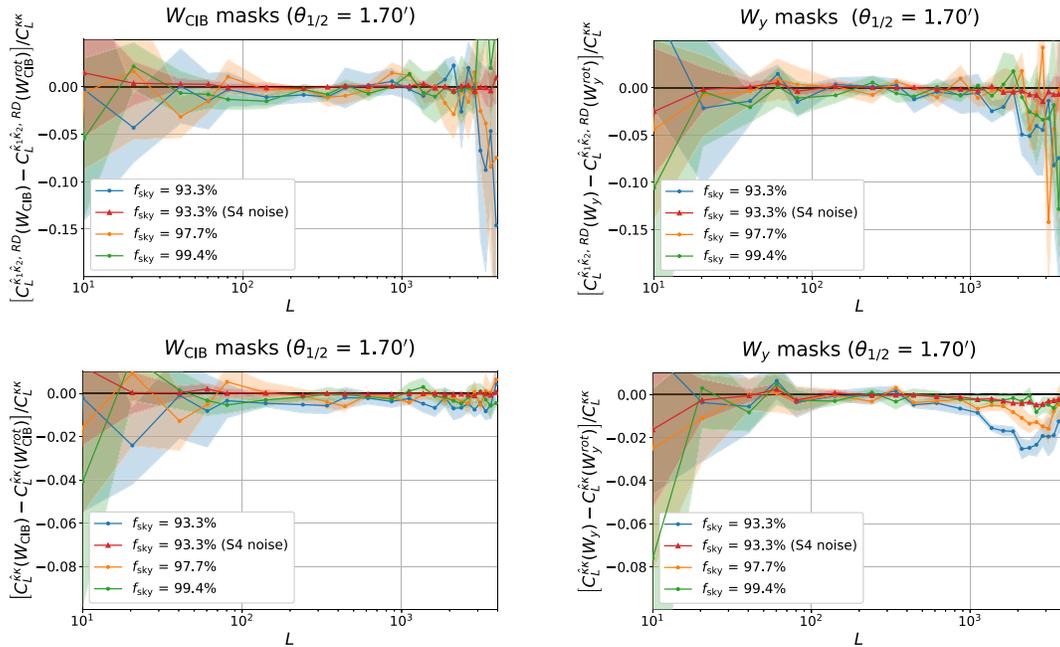


FIG. 11. Bias due to the LSS-correlated masking on the reconstructed CMB lensing convergence autospectrum (top panels) and cross-spectrum (bottom panel) for a subset of the cases shown in Fig. 8. For these results, the masks were constructed smoothing the foreground fields with a $\theta_{1/2} = 1.7'$ Gaussian beam prior to thresholding (instead of $\theta_{1/2} = 5.1'$ for Fig. 8).

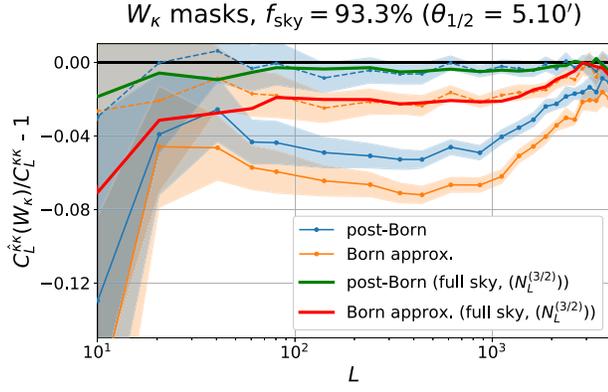


FIG. 12. Differences between the cross-spectra of the reconstructed CMB lensing fields and the true fields relative to the full-sky power spectrum of the true fields. All these curves include mean-field subtraction and f_L^{MC} normalization correction. The results are obtained by performing lensing reconstruction after masking the sky with W_κ masks constructed from the post-Born and Born κ maps of Ref. [67] (solid blue and orange, respectively). Dashed lines show the results obtained using the LSS-uncorrelated masks W_κ^{rot} . These are consistent with the results of lensing reconstruction on full sky (green and red lines), where the observed discrepancy is only due to $N_L^{(3/2)}$.

modifies the shape of the $N_L^{(3/2)}$ bias as shown in Fig. 12, in which the results obtained with rotated masks (dashed lines) are consistent with those obtained on the full-sky (green and red lines, which are consistent with the expected $N_L^{(3/2)}$ reconstruction bias). Our results for correlated mask bias are again calculated as differences between spectra on rotated and unrotated masks. As such, the $N_L^{(3/2)}$ bias is mainly removed when taking this difference, and the impact of post-Born lensing is only important to the extent that the correlated mask changes the post-Born contribution to $N_L^{(3/2)}$ compared to the Born case. Since the amplitude of $N_L^{(3/2)}$ is relatively small, the overall impact of post-Born lensing is small (see Fig. 13). Neglecting post-Born effects is therefore not expected to significantly affect the other results shown for other LSS-correlated masks.

B. Cross-correlation with CMB lensing

In the previous section, we showed that using optimal filtering in CMB lensing reconstruction recovers some of the information lost by masking LSS-correlated sky areas, reducing the expected biases in the reconstructed CMB power spectrum. However, this does not guarantee that the CMB lensing field is properly recovered at the sky location masked in the CMB maps used for lensing reconstruction. In this section, we quantitatively assess this problem using statistics involving cross-correlation between CMB lensing and external tracers.

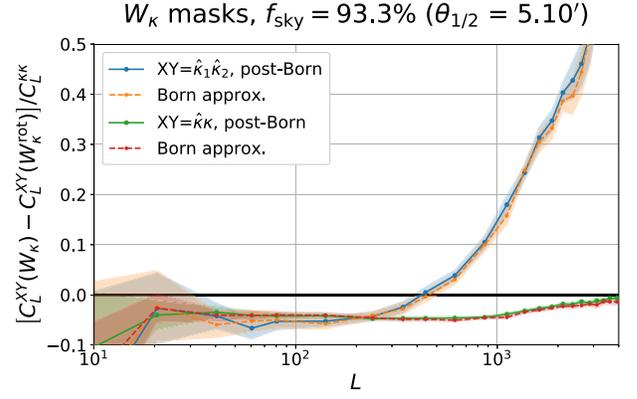


FIG. 13. Bias due to the LSS-correlated masking on the reconstructed CMB lensing convergence autospectrum (in blue and orange) and cross-spectrum (in green and red) relative to the true full-sky $C_L^{\kappa\kappa}$ for W_κ masks. Results including post-Born lensing are shown in solid lines, and those in the Born approximation are shown as dashed lines. All these curves include the f_L^{MC} correction and mean-field subtraction but retain the $\text{RD-}N_L^{(0)}$ bias.

1. CIB and tSZ cross-correlation

To measure the mask biases on $C_L^{\kappa^{\text{CIB}}}$ and $C_L^{\kappa^{\text{tSZ}}}$, we used a pipeline similar to the one we used to measure the biases on $C_L^{\kappa\kappa}$. We computed the cross-correlation power spectrum between the *WEBSKY* y and CIB map at 545 GHz with the lensing map reconstructed from CMB maps masked with our LSS-correlated masks. We then subtracted the same results calculated using the LSS-uncorrelated (rotated) masks. As for the CMB lensing autospectrum, this procedure isolates the effect of the LSS masking from other reconstruction biases, which in this case are only due to the $N_L^{(3/2)}$ bias [18]. We chose in particular the 545 GHz for the CIB emission as it has been successfully employed in delensing studies and offers a good trade-off between signal to noise and dust contamination [72,83,84]. In the left panel of Fig. 14, we show that the $N_L^{(3/2)}$ bias in cross-correlation obtained for the lensing field reconstructed from an LSS-uncorrelated mask and an SO noise level is consistent with the result obtained from full-sky reconstructions. The full-sky $N_L^{(3/2)}$ bias for $C_L^{\kappa^{\text{y}}}$ is consistent with the one expected for a tracer probing the matter density at a median redshift of $z \approx 0.6$ without post-Born corrections (see, e.g., Fig. 13 of Ref. [18]).

The other panels of Fig. 14 show the amplitude of the mask bias relative to the true cross-correlation power spectrum of the *WEBSKY* maps for a few masks representative of masking infrared sources and galaxy clusters for an SO-like survey. For $C_L^{\kappa^{\text{y}}}$, masking infrared sources could bias the measured power spectrum by $\sim 2\%$ for an experiment with a *Planck*-like angular resolution, or half this for an experiment with an SO-like resolution. Cluster masking

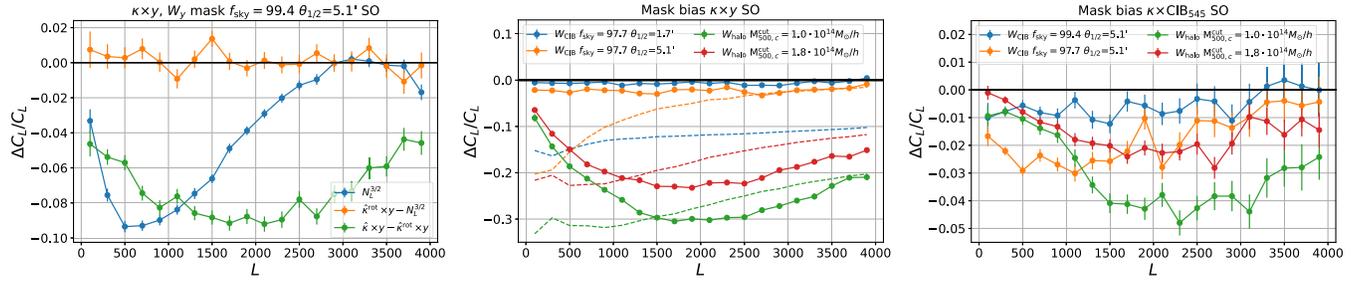


FIG. 14. In the left panel, we show an example of the bias induced by the LSS-correlated masking on the cross-correlation power spectrum between CMB lensing and tSZ Compton y parameter (green) for SO. In blue, we show the $N_L^{(3/2)}$ bias to C_L^{ky} due to the non-Gaussianity of the lensing field itself measured on the full sky. The cross-correlation power spectrum of the reconstructed CMB lensing field with the y map when masking the sky with an LSS-uncorrelated mask is consistent with $N_L^{(3/2)}$ (orange) measured on the full sky. The middle and right panels show the mask bias to the cross-correlation power spectrum between CMB lensing and y and the CIB emission at 545 GHz, respectively, for the most significant masks considered in this work. Such biases are larger than the one obtained for the reconstructed CMB lensing autospectrum and cross-spectrum shown in Figs. 8 and 9. Dashed lines show the expected fractional difference of the power spectrum computed using the true κ fields on a masked sky compared to its expected value on the full sky (for display purposes shown divided by a factor of 2). The optimal filtering used for the lensing reconstruction recovers a large fraction of the information lost by masking.

has a more dramatic impact, introducing biases of 10%–30% in the measured power spectrum, with a broad peak around the median scales of the masked objects (for our case roughly $6'$). This case is of course an extreme example useful to quantify the error on the recovered correlation inside the masked regions; for practical analyses, using C_L^{ky} directly for scientific purposes cluster masking is avoided. However, the optimal filtering recovers a significant amount of information: applying the LSS-correlated mask on the true field would reduce the C_L^{ky} power spectrum by a much larger amount. As an example, in the middle panel of Fig. 14, the dashed lines show the fractional difference between C_L^{ky} computed on a masked sky and the one computed on the full-sky using the true κ field. Even for the most extreme cases of cluster masking, optimal filtering recovers about 80% of the signal for $L \lesssim 500$ and reduces the bias by a factor of ~ 2 at $L \gtrsim 1000$.

For $C_L^{\kappa\text{CIB}}$, despite the non-negligible correlation between CIB and tSZ due to infrared emission of galaxies in galaxy cluster environments [41,85], the mask biases are smaller ($\lesssim 5\%$) for the masks considered for this study. For the most realistic cases related to infrared point sources shown in Fig. 14, we found a reduction of power of about 2% approximately constant for the scales most relevant for delensing ($L \lesssim 2000$). As these variations are small, we do not expect the delensing efficiencies for CIB-delensing to be strongly impacted by masking biases.

We performed a similar analysis with CMB lensing maps reconstructed from CMB maps with an S4-like noise level. We focused on a few specific aggressive masks: W_y and W_{CIB} threshold masks removing $f_{\text{sky}}^{\text{mask}} = 7.7\%$, computed with a Gaussian smoothing of $\theta_{1/2} = 1.7'$ and W_{halo} with $M^{\text{cut}} = 10^{14} M_{\odot}/h$. For the latter, shown in Fig. 14 for SO noise levels, we observed a reduction of the bias on C_L^{ky} and

$C_L^{\kappa\text{CIB}}$ by a factor of ~ 2 for $L \lesssim 2000$ and by a factor of ~ 3 for $2000 \lesssim L \lesssim 4000$ thanks to the improved performances of the optimal filtering. Similar trends can be observed for the threshold masks where the reduction of the biases compared to an SO-like noise is even larger: going from an SO-like to an S4-like noise, for the W_y mask the median bias across all the angular scales changes from 16% to 6% for C_L^{ky} and from 1.6% to 0.2% for $C_L^{\kappa\text{CIB}}$.

2. Cluster mass calibration

The abundance of galaxy clusters as a function of mass and redshift is a highly sensitive probe of cosmology: it strongly depends on the growth rate as well as on the geometry of the universe. The main challenge for the application of cluster abundance in cosmology is the inference of the true mass of the cluster from observable quantities such as X-ray, tSZ luminosity, or optical richness. Gravitational lensing of light sources behind clusters is one of the most promising techniques to estimate their masses as it is sensitive to the total matter distribution and less affected by complex details of baryonic physics in dense environments. CMB-cluster lensing might be particularly useful for estimating the mass of high-redshift clusters, for which it is difficult to observe background galaxies with sufficient sensitivity, and for providing estimates complementary to galaxy weak lensing for low redshift clusters as it is sensitive to different systematic effects [86].

Since the signal-to-noise ratio expected for each cluster in CMB lensing maps, even for futuristic surveys, is well below 1 for clusters of $M_{500,c} \approx 10^{14} M_{\odot}/h$ [87], cluster masses are usually computed as the average mass of a set of clusters (e.g., Refs. [88,89]). We therefore quantified the impact on the recovered mean cluster halo mass by stacking

the CMB lensing maps reconstructed from different masked fields at the location of the clusters. We selected the objects in the WEBSKY halo catalog, mimicking a complete mass-limited sample for SO-like noise described in Sec. III B, and estimated the mass of the clusters from the radial profile of the stacked convergence map with the CMBHALOLENSING public code,¹⁴ For this purpose, we stacked cutouts of $25' \times 25'$ stamps around the location of each cluster from our full-sky maps and binned the radial profile of the stack in 25 radial bins. We estimated the covariance of each radial profile by splitting our cluster sample into 192 subregions corresponding to the objects located within a HEALPIX pixelization of $n_{\text{side}} = 4$ and adopting a jackknife resampling approach following Ref. [90]. In the fit, we assumed the redshift of the stack to be equal to the mean redshift value of the cluster sample $\bar{z} = 0.63$ and used points at a distance $r < 10'$; we also accounted for the mass dependence of the halo profile concentration parameter c using the scaling relation of Ref. [52]. We then quantified the mask bias by comparing the fitted mass value with the one obtained by stacking the full-sky reconstructed κ map. To account for the finite sky coverage of future ground-based surveys such as SO and S4, we rescaled the jackknife covariance estimated over the full sky by the sky fraction covered by those experiments before performing the fit. For this purpose, we assumed a common sky fraction of $f_{\text{sky}}^{\text{obs}} = 50\%$.

Several methods have been proposed to extract the CMB-cluster lensing signal from CMB temperature and polarization maps with different levels of optimality and sensitivity to extragalactic foreground contaminations [17,91–96]. The standard QE in particular is known to be suboptimal and biased-low for low CMB noise levels [95,97,98]. We quantified this bias by computing the mean cluster mass from the radial profile obtained by stacking the WEBSKY κ at the cluster locations and assuming the same covariance of the full-sky reconstructed map. This value is consistent with the mean mass of the sample estimated from the values present in the WEBSKY halo catalog ($\bar{M}_{500,c} = 2.64 \times 10^{14} M_{\odot}/h$), while the one obtained with the full-sky reconstructed map is biased low by $\approx 1\sigma$ ($\hat{M}_{500,c}^{\text{full-sky}} = (2.29 \pm 0.15) \cdot 10^{14} M_{\odot}/h$). As a consistency check, we estimated the mean mass from the profile obtained from convergence maps reconstructed from CMB maps masked with W_X^{rot} and found it to be consistent with the results obtained from the full-sky reconstructed maps.

Figure 15 shows a comparison of the radial profiles obtained from different LSS-correlated masks. The changes in the halo profile are mainly concentrated toward the halo center, while the largest scales are mainly unaffected. The CIB masks that remove the brightest infrared sources induce changes at the $\approx 1\%$

level in the profile core but do not significantly affect the mass estimation as the differences with results obtained using the full-sky reconstructed mask are $\lesssim 0.5\sigma$. Masking all the clusters in the sample before lensing reconstruction shows that only about 30% of the signal at the halo location can be recovered as the recovered mass is $\hat{M}_{500,c} = (0.84 \pm 0.17) \cdot 10^{14} M_{\odot}/h$. As discussed above, this masking choice is overconservative; however, results obtained with W_{halo} and W_y masks are similar if the removed sky fraction is comparable, as these masks are correlated.

We therefore also considered a less extreme case where just the brightest clusters are masked, computing the cluster mass from a convergence map reconstructed from a CMB map masked with a W_y mask removing $f_{\text{sky}}^{\text{mask}} = 2.3\%$ at a smoothing scale of $1.7'$. This removes a similar sky fraction to a halo mask with the redshift-dependent selection function of the *Planck* cluster catalog ($f_{\text{sky}}^{\text{mask}} = 1.6\%$) [99] if the clusters are masked within a radius of $2\theta_{500,c}$ from their centres. In this case, we observe a bias on the recovered mean cluster mass of about $\approx 2.5\sigma$. To investigate if the mask bias affects clusters in particular redshift bins, we repeated the analysis, masking only clusters at $z < 0.6$, $z > 0.6$ and $z > 1$. Accounting for the increased statistical uncertainties due to the lower number of sources, we found that the bias is reduced to 1.1σ for clusters at $z > 0.6$, while for objects at $z > 1$, the recovered mass is consistent with the one expected from the sample within $\sim 0.5\sigma$. The bias, however, becomes much more important for low-redshift clusters, where for objects at $z < 0.6$ we detected the mask bias at $\sim 3.5\sigma$. This is likely due to the fact that the size of the masked object is about twice as large then the one at $z > 0.6$ and therefore optimal filtering is less effective in recovering information. The mask bias is still measurable at $\sim 1.5\sigma$ significance for less aggressive W_y masks removing $f_{\text{sky}}^{\text{mask}} = 0.6\%$ at a smoothing scale of $1.7'$. These results suggest that the mask biases will not significantly affect the science case related to cluster mass calibration from CMB lensing mass in the regime where this is the most accurate technique, but more care has to be taken when calibrating the mass of low redshift objects.

We repeated the analysis on CMB lensing maps reconstructed from CMB maps having an S4-like noise level. We focus in particular on a very aggressive case where we masked the sky prior to reconstruction with a W_{halo} mask with $M^{\text{cut}} = 10^{14} M_{\odot}/h$ (consistent with an S4-like cluster sample), and we later stacked the reconstructed κ at the location of the same halos that were masked. Despite being unrealistic, such an extreme case is useful to assess the performances of the optimal filtering in recovering small-scale information in the maps. We found that the estimated mass from the stack has a bias at $\sim 1.2\sigma$ level compared to the results obtained stacking a κ map reconstructed from full sky observations. Optimal filtering recovers about 90%

¹⁴<https://github.com/simonsobs/cmbhalolensing>.

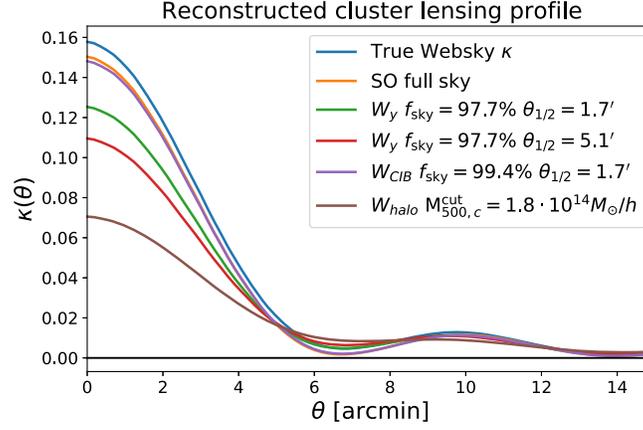


FIG. 15. The azimuthally averaged radial profile of the convergence maps stacked at the location of SO detected clusters. Different colors show the results obtained on κ maps reconstructed from different masked CMB maps and assuming an SO noise level in CMB maps. Masking CMB maps with LSS-correlated masks causes a suppression of the signal at small scales compared to the full-sky reconstruction and modifications of the profile at intermediate scales when masks remove large sky fractions. The full-sky results do not retrieve an unbiased profile as they have been derived with a standard QE reconstruction; as such, results obtained stacking directly on the CMB lensing convergence maps of the WEBSKY suite are shown for reference to quantify the importance of the QE bias.

of the expected signal at the halo location for S4 noise levels. The same estimate assuming SO-like noise levels and the same mask gives a much larger bias ($\sim 7\sigma$) in the estimated cluster mass with only about 50% of the expected signal properly recovered. Furthermore, we found that the conclusions on the redshift dependency previously discussed for the SO noise level and the SO cluster sample hold also for the S4 noise and cluster sample, with a mask bias increased to $\sim 3\sigma$, $\sim 1.8\sigma$, and $\sim 1\sigma$, for objects located at $z < 0.6$, $z > 0.6$, and $z > 1$ respectively. Since the mask bias we considered here for S4 gives a marginal detection significance, despite being an extreme xoverconservative case, we do not expect realistic foreground masks to severely impact cluster mass calibration for S4.

Given that the estimator we employed in this section is not only biased but suboptimal, and given that our noise level is higher than a full minimum-variance lensing reconstruction (we only used temperature data), a bias in the estimated cluster mass over the full sample could be important even if only the brightest fraction of the detected clusters is conservatively masked prior to the reconstruction. However, a more targeted analysis should be carried out to accurately assess this effect for future datasets including all the complexity of the cluster selection function of each experiment.

3. Mask-CMB lensing deflection correlation

In Paper I, we presented an analytic model of the correlated mask bias on standard CMB pseudo- C_ℓ angular power spectrum estimators. This can be evaluated as an effective correction to the lensed CMB correlation function $\tilde{\xi}(r)$ given by

$$\Delta\tilde{\xi} \approx \partial_r \tilde{\xi}(r) \bar{\Delta}(r) \quad (5.1)$$

and then converted into a correction on CMB C_ℓ 's. In the last equation, $\bar{\Delta}(r)$ is the average over the unmasked area of the change in the separation of points due to lensing,

$$\bar{\Delta}(r) = 2 \frac{\langle \alpha_r(\mathbf{x}) W(\mathbf{x}) W(\mathbf{x}') \rangle}{\langle W(\mathbf{x}) W(\mathbf{x}') \rangle}, \quad (5.2)$$

where \mathbf{x} , $\mathbf{x}' = \mathbf{x} + \mathbf{r}$ are directions in the sky, \mathbf{r} is their separation vector, and α_r is the component of the deflection field $\boldsymbol{\alpha} = \nabla\phi$ parallel to \mathbf{r} . $\bar{\Delta}(r)$ does not have a general analytical expression but can in principle be calculated empirically. The numerator of Eq. (5.2) can be expressed in terms of the cross-correlation power spectrum between the E-modes of the spin-1 masked deflection field αW , denoted by E , and the sky mask W . In the flat sky, this reads

$$\langle \alpha_r(\mathbf{x}) W(\mathbf{x}) W(\mathbf{x}') \rangle = - \int \frac{dl}{2\pi} l C_l^{\text{EW}} J_1(lr). \quad (5.3)$$

Thus, if we have an estimate of C_L^{EW} , we can calculate the CMB power spectrum bias for any mask. This can be constructed from simulations where we know κ (and hence $\boldsymbol{\alpha}$) and the mask W . Alternatively, this can be estimated from data if we have reliable estimates of the lensing deflection field. We used this technique in Paper I to evaluate the bias induced by a threshold mask built on *Planck* GNILC maps on the CMB temperature power spectrum of *Planck* SMICA maps and showed that our analytical prediction so computed matched the mask bias observed on data.

The estimator of the CMB C_ℓ bias using C_L^{EW} requires being able to estimate the deflection field (to calculate C_L^{EW}). If a masked lensing reconstruction is used to estimate this, the estimate of C_L^{EW} may itself be biased due to the bias

on the lensing reconstruction. To test this, we converted the reconstructed lensing potential $\hat{\phi}$ into a deflection field $\hat{\alpha}$ with a spin-1 spherical harmonics transform.¹⁵ We then computed the cross-correlation power spectrum $C_L^{\hat{E}W}$ between the reconstructed masked deflection field $\hat{\alpha}W$ and W and compared it with C_L^{EW} directly computed using the deflection field extracted from the WEBSKY κ map. We focused in particular on the most important realistic cases analyzed in Paper I, i.e., the W_{CIB} and W_y masks with $f_{\text{sky}} = 99.4\%$ and for both $\theta_{1/2} = 5.1'$ and $1.7'$, which could all be detectable at more than 5σ significance for SO noise level. We found the error on C_L^{EW} for SO noise levels to be lower than 30% for $L \lesssim 3000$ and lower than 10% for the CIB masking which cannot be avoided by component separation. We evaluated Eq. (5.3) with $C_L^{\hat{E}W}$ and found that the analytically predicted mask bias on C_ℓ^{TT} is consistent with the one computed using the true C_L^{EW} measured from WEBSKY to better than $\sim 0.5\%$ on $L \lesssim 4000$. Even for the most aggressive masks considered here, i.e., W_{CIB} and W_y that remove $f_{\text{sky}}^{\text{mask}} = 6.7\%$, the error introduced in the prediction of the mask bias using $C_L^{\hat{E}W}$ is less than 10%. For S4 noise levels and the same aggressive masks, the errors on C_L^{EW} are reduced up to a factor of 2 on scales $L \gtrsim 1000$ compared to the SO noise case, and therefore we expect the accuracy of the analytical predictions of the biases to further improve at lower experimental noise levels.

The lensing reconstruction mask bias therefore appears to be sufficiently small that the CMB C_ℓ bias can be estimated accurately enough to reduce its statistical significance below the detection level.

VI. FORECASTS FOR FUTURE CMB EXPERIMENTS

In the previous sections, we have shown that LSS-correlated masks can introduce biases on the reconstructed CMB lensing power spectrum that could become non-negligible if not accounted for. We have shown that the biases on the CMB lensing power spectrum are likely to be negligible for radio sources for the immediate future but may be more important for other masks (a $\sim 2\%$ effect, which can increase to $\sim 5\%$, though, on scales where the lensing reconstruction noise is important; see Figs. 8 and 11).

Below, we estimate the detectability of such mask biases for SO and S4 in terms of cumulative signal to noise where we used as the signal the mask biases measured in Sec. V and the noise includes the sample and noise variance of each experiment. For this purpose, we assumed a sky coverage of $f_{\text{sky}}^{\text{obs}} = 40\%$ and the realistic publicly available noise power spectra for the lensing potential reconstructed

using only temperature modes.¹⁶ We fix all cosmological parameters, so the results are an upper limit on the impact on any cosmology constraint.

The results for all the W_X masks considered in this analysis are summarized in Fig. 16, in which we show the detection significance of the mask biases as a function of maximum lensing multipole L_{max} included in the analysis. We do not include the results for the W_κ mask since κ is not directly observable, but for comparison, we also show the expected biases from directly masking the κ convergence field as done in Sec. III C. Our results clearly show that, using optimal filtering of the CMB maps for the lensing reconstruction, the biases are much smaller than they would be if κ were masked directly (which would give a signal detectable with high significance, well above 5σ in future measurements).

For W_{halo} , masking up to 1%–2% of the sky (consistent with the sky fraction removed by masking SO-like tSZ cluster sample), the mask bias can be measured with a statistical significance $\lesssim 2\sigma$, while for cluster samples more similar to the S4 ones, where the mass limit of $10^{14} M_\odot/h$ and the masked sky fraction is $\sim 10\%$, the mask bias can be detected at $\lesssim 3\sigma$ in the reconstructed lensing field, using both the SO and the S4 noise. Similar results can be observed for the threshold mask W_y when the foreground field has been smoothed with $\theta_{1/2} = 5.1'$ Gaussian beam and a similar sky fraction as in the W_{halo} case is masked, given that the two types of masks are correlated.

We note that cluster masking may not be used in final CMB lensing analyses, and the tSZ contamination can also be reduced by analyzing component-separated CMB maps where components with a tSZ-like SED (spectral energy distribution) are projected out during component separation or using dedicated modified quadratic estimators [17,100]. To assess the validity of the methods and to quantify potential residual emission in these alternative analyses, however, it is common practice to compare the results with more conservative analyses based on cluster masking as also done in the *Planck* 2018 lensing analysis [21]. As such, our results show that care is required when performing these kind of comparisons, as mask biases may lead to misleading inconsistencies in the comparison of the CMB lensing estimates.

For the W_{CIB} (with $\theta_{1/2} = 5.1'$) masks, the bias detectability remains $\lesssim 1\sigma$ when the masked sky regions removes $\sim 5\%$ of the sky or less. For more aggressive masks, removing $f_{\text{sky}}^{\text{mask}} \approx 7\%$, the bias will be detectable at more than 3σ significance. The increased significance is not unexpected since the CIB is highly correlated with the CMB lensing field, especially at $L \lesssim 1000$, and therefore a

¹⁵For this purpose we used the relation between the lensing potential and the E-modes of the deflection field $\hat{\alpha}_{LM}^E = \sqrt{L(L+1)}\hat{\phi}_{LM}$ and assumed $\hat{\alpha}_{LM}^B = 0$.

¹⁶For SO, we used the so-called baseline noise from https://github.com/simonsobs/so_noise_models. For S4, we took the noise power spectra available at https://cmb-s4.uchicago.edu/wiki/index.php/Survey_Performance_Expectations.

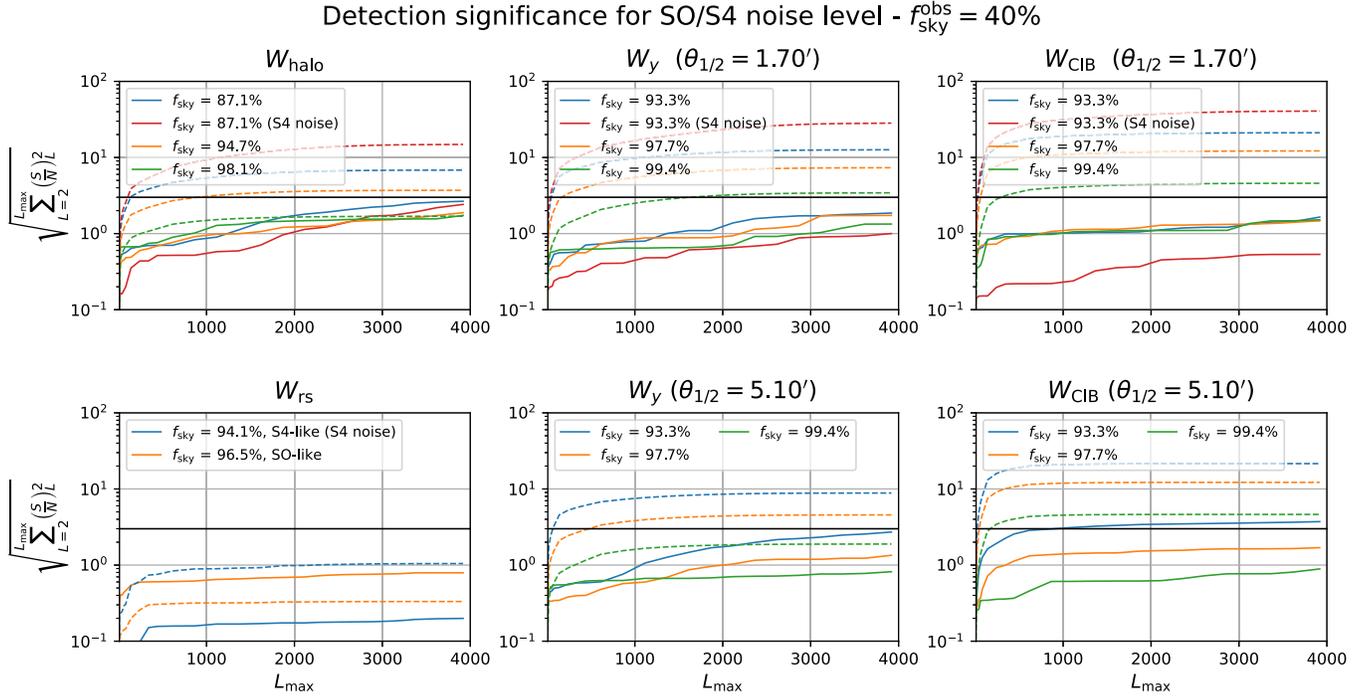


FIG. 16. Detection significance of the mask bias in the CMB lensing power spectrum for future high-resolution ground-based experiments as a function of maximum multipole L_{max} included in the analysis and for different foreground masks. We considered an observed sky fraction $f_{\text{sky}}^{\text{obs}} = 40\%$. We assumed an SO-like lensing noise for all cases except for those shown in red which assumed an S4-like CMB lensing noise. The dashed lines show the bias obtained directly masking the κ field (see Sec. III C), while the solid lines show the bias on the reconstructed κ field from masked CMB maps (see Sec. IV). The horizontal black line highlights a statistical detection significance of 3σ for reference.

mask that removes these peaks naturally enhances the mask bias effect.

The detection significance for both W_y and W_{CIB} masks built from higher-resolution observations (i.e., with $\theta_{1/2} = 1.7'$ smoothing) will produce biases that will be measured with lower statistical significance compared to the $\theta_{1/2} = 5.1'$ cases. They always stay $\lesssim 2\sigma$ for SO and $\lesssim 1\sigma$ for the S4 cases we considered, showing the improved performance of the optimal filtering in recovering information in presence of lower noise level and smaller holes.

For radio sources, W_{rs} , the detection significance of the bias is below the 1σ level, whether we are masking all the radio sources up to the detection limit of SO or S4.

We finally stress that the bias for any Poisson source mask is expected to be very small, as for radio sources, since the mask is largely determined by random Poisson sampling of the background distributions rather than tracing lensing-correlated perturbations closely. For example, as long as masks removing infrared sources are only removing Poisson sources (i.e., at very high or low redshift), rather than peaks of the full CIB field, their bias should also be negligible. Although masking peaks in the CIB emission is not common practice in CMB analysis, our W_{CIB} masks can potentially mimic the effect of masking bright infrared point sources, where a fraction of them is from strongly lensed objects that are therefore highly

correlated with the matter distribution along the line of sight on much smaller scales than CMB lensing is sensitive to [101]. We note that WEBSKY simulations were not constructed to reproduce the source number counts in the infrared nor include any effect of magnification bias which would affect the number of detected sources in CMB maps for a fixed noise level. However, preliminary analyses showed a good agreement between the expected source number counts at the highest fluxes from semianalytical models [102,103] and the source counts expected from the halo-model adopted in the WEBSKY simulation to construct the CIB maps [104]. Our WEBSKY simulation results should therefore provide a reasonable ballpark estimate of the total effect, since the uncorrelated Poisson part of the distribution only affects the mask bias via the total f_{sky} masked.

We performed a similar forecast analysis for the mask biases in the cross-correlation between CMB lensing and external tracers presented in Sec. V B 1. In this case, we assumed the publicly available SO and S4 noise power spectra for the y map achievable with an ILC (internal linear combination) component separation and for the lensing noise; for the CIB, we fitted a white noise level and an effective beam from the noise-dominated regime of the beam-deconvolved power spectrum of the GNILC maps at 545 GHz. We obtained a noise level of 4.84 Jy/sterad (janskys per solid angle) and a Gaussian

beam with $\theta_{1/2} = 4.65'$. We also considered future CIB measurements of CCAT-prime and assumed the beam and noise power spectra described in Ref. [105]. As expected, biases in cross-correlation are more important than those reported above for the CMB lensing spectrum. We considered mainly threshold masks smoothed with $\theta_{1/2} = 1.7'$, as it is the case more relevant for future experiments and W_{halo} masks. For SO, we found that W_{CIB} and W_y threshold masks do not produce any significant bias on $C_L^{\kappa\text{CIB}}$ if they remove less than 6% of the sky both for *Planck* and CCAT-prime noise levels and resolution. Biases due to W_y masks are, however, more harmful for the analysis of $C_L^{\kappa y}$ even if only a minor fraction of the brightest clusters is removed. W_y masks biases will in fact be detectable with a significance of $\sim 3\sigma$ if the mask removes 0.6% of the observed sky and of 5σ when masking 2.3% of the sky (consistent with the sky fraction removed by masking *Planck* tSZ-detected clusters).

The significance increases to 10σ and 13σ for W_{halo} masks removing SO- and S4-like clusters, respectively. These masks will also leave detectable biases in $C_L^{\kappa\text{CIB}}$ at 2.5σ and 4.5σ , respectively. For an S4-like survey, we saw in Sec. V B 1 that the mask biases are reduced compared to an SO-like one; however, the noise of the reconstructed CMB lensing and y maps is also lower, and the detection significance might still be comparable to the ones of an SO-like survey. Nevertheless, we found that the overall detection significance of the mask biases of a given mask is reduced for an S4 survey compared to an SO-like one. Even for the extreme cases discussed in Sec. V B 1, the mask biases on $C_L^{\kappa\text{CIB}}$ are negligible as they will be measured with a detection significance $\lesssim 1.5\sigma$. For $C_L^{\kappa y}$, biases induced by W_{CIB} will not be detectable, and the significance of mask biases induced by W_y and W_{halo} masks will be reduced by about a factor of 2. This reduction is enough to reduce the detection significance of mask biases of *Planck*-like detected clusters to marginal significance ($\sim 2.5\sigma$), but more aggressive cluster masking will introduce biases in $C_L^{\kappa y}$ that will be detectable at high significance.

Finally, we note that if CMB polarization is also used in the lensing reconstruction, and the polarization mask is much smaller than the temperature mask as expected, all the biases that we found here for the temperature are expected to be reduced. Also, combined minimum-variance and optimized estimates [106–108] of the lensing potential should have these biases reduced as the relative importance of polarization-based reconstruction channels will increase for future observations.

VII. CONCLUSIONS

In this work, we studied the impact that foreground masks correlated with the large-scale structure distribution could have on the reconstructed CMB lensing poten-

tial map and power spectrum. Future high-resolution ground-based CMB experiments, such as SO and S4, will resolve much larger populations of extragalactic sources than current experiments, so masking larger areas of resolved tSZ-selected clusters and radio sources may be necessary and its impact must be carefully quantified. In Paper I, we already showed how such masks can potentially give large biases on the CMB temperature and polarization power spectra, even if the masked sky area is small. Building on these results, in this paper, we have shown that:

- (i) Extragalactic foregrounds masks that are correlated to CMB lensing give substantial biases on simple lensing power spectrum estimates if the lensing convergence field were masked directly (or no information inside the masked sky area can be recovered).
- (ii) Significantly smaller biases are obtained on the reconstructed CMB lensing field and power spectrum if the CMB fields used for the reconstruction are optimally filtered, effectively recovering signal inside small mask holes.
- (iii) The simple halo model or analytic Gaussian models that we derive can provide a qualitative understanding of the effect when the lensing field is masked directly, though non-Gaussianity of the extragalactic foregrounds is important to model for accurate results.
- (iv) Radio source masks, and any other mask constructed from Poisson sources, should give a bias that is safely negligible. Masking foreground peaks in the tSZ or the infrared emission, or large halos, can instead give larger biases and should be avoided if possible.
- (v) Biases on the reconstructed CMB lensing power spectrum from CMB temperature will mostly be measured with only marginal significance ($\lesssim 2\sigma$) for forthcoming experiments.
- (vi) Biases in the cross-correlation between CMB lensing and tSZ or CIB are detected with higher significance for realistic foreground masks ($\lesssim 3\text{--}4\sigma$ for SO) but can become much larger for more aggressive masks. Thanks to improved performance of the optimal filtering in presence of low noise levels, the importance of such biases is greatly reduced for S4, however, not enough to make them completely negligible.
- (vii) At SO noise levels, cluster mass calibration with CMB lensing can be significantly affected by mask biases only for low- z objects but will still deliver unbiased results at high- z . This holds even if only the brightest clusters are masked. Cluster mass calibration for S4, conversely, will likely not be affected by mask biases when realistic foreground masks are employed.

(viii) Since masking biases in the reconstructed CMB lensing are small, the large masking biases expected for CMB power spectra can be predicted and marginalized self-consistently using CMB lensing as an LSS tracer and the analytical model we described in Paper I.

We showed that mask biases are small for lensing reconstruction using optimal CMB filtering, but alternative methods based on sky map inpainting or source subtraction at the map-making level, combined with an isotropic QE instead of optimal filtering, might also be effective in avoiding sharp mask edge effects and recovering information inside small holes, but would have to be assessed separately. Applying an additional optimal filtering step on the reconstructed lensing field may also be able to recover information in small mask holes that were not filled at the level of the CMB map (cf., Ref. [76]).

A detailed analysis including CMB polarization in the lensing reconstruction as well as polarized foregrounds is left for future work, but biases are likely to be smaller given the expected levels of polarization of radio and infrared sources and other extragalactic emissions. A study of the impact of these correlated mask biases on the estimation of cosmological parameters is also left for future works.

ACKNOWLEDGMENTS

We thank Anthony Challinor for useful discussion and Boris Bolliet, Emmanuel Schaan, and Blake Sherwin for comments. A.L., G.F., and J.C. acknowledge support from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013)/ERC Grant No. 616170 and support by the UK STFC Grants No. ST/P000525/1 (A.L.) and No. ST/T000473/1 (A.L. and G.F.). G.F. acknowledges the support of the European Research Council under the Marie Skłodowska Curie actions through the Individual Global Fellowship No. 892401 PiCOGAMBAS. J.C. acknowledges support from a SNSF Eccellenza Professorial Fellowship (Grant No. 186879). M.L. acknowledges the support of the Fondazione Angelo Della Riccia fellowship. The sky simulations used in this paper were developed by the WEBSKY Extragalactic CMB Mocks team, with the continuous support of the Canadian Institute for Theoretical Astrophysics (CITA), the Canadian Institute for Advanced Research (CIFAR), and the Natural Sciences and Engineering Council of Canada (NSERC), and were generated on the Niagara supercomputer at the SciNet HPC Consortium [109]. SciNet is funded by the Canada Foundation for Innovation under the auspices of Compute Canada, the Government of Ontario, Ontario Research Fund—Research Excellence, and the University of Toronto.

APPENDIX A: CMB LENSING RECONSTRUCTION RESPONSES FOR MASKS WITH HOLES

Masking the CMB maps impacts the lensing reconstruction even for a mask that is independent of large-scale structure. The CMB over the masked area cannot be perfectly recovered by optimal filtering, and after mean-field subtraction, a signature remains in the estimators that is inaccurately captured by the standard isotropic lensing responses [73] devised for a full-sky analysis. In CMB lensing estimation analyses, various corrections including this one are usually encapsulated inside a single Monte Carlo correction, at the map or power spectrum level [16,72,110].

This masking effect cannot be perfectly separated on data from correlated-mask signatures and the bispectrum contribution $N_L^{(3/2)}$, but simulations with the same mask can be used to estimate the effect by analyzing results from Gaussian realizations. In this appendix, we aim at some quantitative understanding of the impact of uncorrelated masking of CMB maps with analytical methods, assuming Gaussian unlensed fields, Gaussian lensing potential, and the perfectly linear response of the QE to lensing.

In order to do this, we calculate the response $\mathcal{R}_{LL}^{MM'}$ (now a dense matrix instead of the isotropic response \mathcal{R}_L) that connects the estimator to the underlying deflection in a number of cases relevant to our study. For lensing, there are two effects: (1) the lensing gradient response gets modified, and (2) masking introduces leakage between the lensing gradient and lensing curl, where the lensing curl estimator picks up a term proportional to ϕ and vice versa.

We use notation consistent to the *Planck* lensing curved-sky pipeline [111], where a separable QE \hat{q} of spin $s_o + t_o$ (1, for lensing) can be written in real space (possibly as a sum of terms) like

$${}_{s_o+t_o}\hat{q}(\hat{n}) = \left(\sum_{\ell m} w_{\ell}^{s_o s_i} {}_{s_i}\bar{X}_{\ell m s_o} Y_{\ell m}(\hat{n}) \right) \cdot \left(\sum_{\ell m} w_{\ell}^{t_o t_i} {}_{t_i}\bar{X}_{\ell m t_o} Y_{\ell m}(\hat{n}) \right), \quad (\text{A1})$$

for a set of in-spins s_i, t_i ; out-spins s_o, t_o ; and weight functions w (left and right leg weights being distinguished from the context). On a masked sky, the CMB \bar{X} fed into the QE is anisotropically related to the true CMB via a filtering matrix F ,

$${}_{s}\bar{X}_{\ell m} = {}_{ss'}F_{\ell\ell'}^{mm'} {}_{s'}X_{\ell'm'}^{\text{cmb}}. \quad (\text{A2})$$

In this work, we use exclusively optimal filtering, which recaptures CMB signal in the CMB maps by reconstructing the maximum *a posteriori* full-sky CMB map, including at locations inside the holes. To do this, we assume a fiducial

likelihood model for the CMB, with fiducial spectra C_ℓ^{fid} , a noise variance map N , and a fiducial beam and transfer function \mathcal{B}^{fid} . The filtering matrix is then

$$F = C^{\text{fid},-1} [C^{\text{fid},-1} + \mathcal{B}^{\dagger,\text{fid}} N^{-1} \mathcal{B}^{\text{fid}}]^{-1} \mathcal{B}^{\dagger,\text{fid}} N^{-1} \mathcal{B}^{\text{sky}}, \quad (\text{A3})$$

where \mathcal{B}^{sky} is the sky beam and transfer function. The matrix is in practice rectangular, since typically reconstruction of only a limited set of sky modes is attempted. Of special interest for us is the ideal case of temperature reconstruction using an azimuthally symmetric boolean mask $W(\cos\theta)$ with homogeneous uncorrelated noise variance N_{lev} and isotropic beam and transfer functions b_ℓ . Writing $Y_{\ell m}(\theta, \phi) = e^{im\phi} \lambda_{\ell m}(\cos\theta)$, the inverse noise matrices in (A2) then become

$$[\mathcal{B}N^{-1}\mathcal{B}]_{\ell\ell'}^{mm'} = \delta_{mm'} \frac{b_\ell b_{\ell'}}{N_{\text{lev}}} A_{\ell\ell'}^m,$$

$$A_{\ell\ell'}^m \equiv 2\pi \int_{-1}^1 d\cos\theta W(\cos\theta) \lambda_{\ell m}(\cos\theta) \lambda_{\ell' m}(\cos\theta).$$

The integral may be evaluated quickly for all ℓ and ℓ' using the recursion relations of the associated Legendre polynomials, as described by Ref. [112].

It is interesting to compare optimal filtering results with the case where the CMB inside the mask holes is simply taken to be zero, so the harmonic modes of the CMB are taken directly on the (possibly apodized-)masked CMB map. In this case, the filtering matrix is

$$F_{\ell\ell'}^{(m,\text{apo})} = \frac{1}{C_\ell^{\text{fid}} + N_{\text{lev}}/b_\ell} \frac{1}{b_{\ell'}} A_{\ell\ell'}^m b_{\ell'}^{\text{sky}}, \quad (\text{A4})$$

where A is the exact same matrix as above where now $W(\cos\theta)$ can include some radial apodization function. We first discuss the simple but illuminating case of a small hole in Appendix A 1. The general case is then dealt with in Appendix A 2.

1. Small hole

In the case of a very tiny hole, large-scale temperature modes inside the hole can be well reconstructed, while only the noise level of the map prevents, to some extent, reconstruction of the small-scale modes. Hence, for a fixed experimental configuration, one expects the correction to the lensing cross-spectrum and autospectrum to scale with the excised area f_{exc} , and to grow with the ℓ_{max} of the CMB fed into the QE (while ℓ_{max} is not noise dominated) but to remain substantially smaller than f_{exc} itself.

In Fig. 17, we show the cross-spectrum correction for our baseline configuration, where $\ell_{\text{max}} = 4000$, $N_{\text{lev}}^T = (6.67 \mu\text{K})^2$ and a $\theta_{1/2} = 1.5'$ beam. The black lines show the total correction

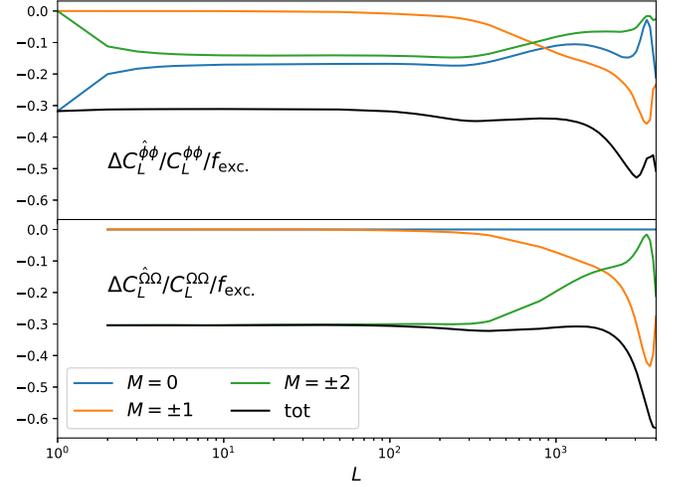


FIG. 17. The figure shows the relative correction to the cross-spectrum of the lensing temperature QE to the true signal due to a mask built out of a population of very small holes, for the baseline experimental configuration in this work. In this regime (for hole size $1/\ell_{\text{max}}$), for fixed ℓ_{max} used in the QE construction and fixed instrumental noise level, the correction scales proportionally to the excised sky area f_{exc} . The lines show the correction divided by f_{exc} , so this scaling is taken out. The top and bottom panels show the gradient and curl part of the lensing QE. The total correction is negative, and significantly closer to the zero value than the -1 value as would be expected if the information were lost proportionally to the masked sky fraction f_{exc} (but not zero, since not all relevant small-scale modes can be accurately recovered at this noise level). The colored lines show the first few multipolar contributions (for a hole centred on the pole), displaying the expected mostly monopolar ($\hat{\phi}$) and quadrupolar ($\hat{\Omega}$) response of the estimator.

$$\frac{\Delta C_L^{\hat{q}q^{\text{in}}}}{C_L^{q^{\text{in}}q^{\text{in}}}} = \frac{1}{2L+1} \sum_{M=-L}^L \frac{\mathcal{R}_{LM,LM}^{qq}}{\mathcal{R}_{LM,LM}^{qq,\text{fid}}}, \quad (\text{A5})$$

in units of the excised area. In the top and bottom plots, \hat{q} stands for the lensing gradient $\hat{\phi}$ and lensing curl potential estimator $\hat{\Omega}$, respectively. The temperature QE gets its information from very small scales, and in our configuration, we still find a number of relevant modes which are too noisy to be reconstructed. This results in a relative correction of $\sim -0.3f_{\text{exc}}$, where the prefactor 0.3 is still somewhat comparable to unity. By comparison, restraining the QE to use $\ell_{\text{max}} = 3000$ or 2000 for the same noise level reduces the large scale correction by a factor of 3 or 40, to $\sim 0.1f_{\text{exc}}$ or $\sim 0.0075f_{\text{exc}}$, respectively.

After rotating the sphere so that the mask covers the pole and is azimuthally symmetric, the response matrix is diagonal in M , and M corresponds for large-scale lenses to the local multipolar structure. The colored lines in Fig. 17 show the first few M multipole contributions. As might be expected, we find a large-scale mostly monopolar

and quadrupolar responses for $\hat{\phi}$, corresponding to distortions in the local estimated spin-0 convergence and spin-2 shear (for $L \geq 2$), respectively. In contrast, there is no monopolar curl potential response, since close to the pole this corresponds to an undetectable rigid rotation. Since the lensing deflections α are of the order of a couple of arc minutes, the distortion at the pole should not affect the reconstruction much farther away than that. Consistent with this, we find that the correction converges very quickly, requiring M_{\max} at most $\sin \alpha \cdot (2L_{\max} + 1) \approx 5$.

2. Larger holes

We now obtain the analytic expression for the response and discuss more general implementations.

For a generic separable QE in Eq. (A1), the response matrix always involves the product of two Gaunt integrals each with one lensing L, M and two CMB ℓm 's. The first

integral contains the observed deflection and two data legs multipoles, together with the QE weights. The second integral contains instead the CMB sky lensing and CMB, together with the response weights. The filtering matrix couples the CMB ℓ 's and m 's together. Concretely, let ${}_{ab}\mathcal{R}$ denotes the linear response of the spin a estimator to the sky spin- b source,

$${}_a\hat{q}_{LM} = {}_{ab}\mathcal{R}_{LL'}^{MM'} {}_b q_{L'M'}. \quad (\text{A6})$$

Also, let $R^{b,s}$ be the response weights of the spin s CMB to the sky spin- b source (for lensing, $R^{\pm 1,s} = \pm \sqrt{(\ell \pm s)(\ell \mp s + 1)}$) [111,113]), and

$$C_{\ell}^{st} \delta_{mm'} \delta_{\ell\ell'} = \langle {}_s X_{\ell m}^{\text{cmb}} X_{\ell' m'}^{\dagger, \text{cmb}} \rangle. \quad (\text{A7})$$

The response matrix takes then the following form,

$${}_{ab}\mathcal{R}_{LL'}^{MM'} = (-1)^a \begin{pmatrix} w^{s_o s_i} & w^{t_o t_i} & 1 \\ \ell_1 & \ell_2 & L \\ m_1 & m_2 & -M \\ s_o & t_o & -a \end{pmatrix} {}_{s_i s} F_{\ell_1 \ell_1' t_i t'}^{m_1 m_1' t_1 t_1'} F_{\ell_2 \ell_2'}^{m_2 m_2' t_2 t_2'}$$

$$\cdot \left\{ \begin{pmatrix} 1 & R^{b,t} C^{t-s} & 1 \\ \ell'_1 & \ell'_2 & L' \\ m'_1 & m'_2 & -M' \\ s & b-s & -b \end{pmatrix}^{\dagger} + \begin{pmatrix} R^{b,s} C^{s-t} & 1 & 1 \\ \ell_1 & \ell_2 & L \\ m_1 & m_2 & -M \\ b-t & t & -b \end{pmatrix}^{\dagger} \right\},$$

where the matrices stand for the weighted Gaunt integrals

$$\begin{pmatrix} c^1 & c^2 & c^3 \\ \ell_1 & \ell_2 & \ell_3 \\ m_1 & m_2 & m_3 \\ s_1 & s_2 & s_3 \end{pmatrix} \equiv \int d^2 n \prod_{i=1}^3 (c_{\ell_i s_i}^i Y_{\ell_i m_i}(\hat{n})), \quad (\text{A8})$$

where always $s_1 + s_2 + s_3 = 0$. In terms of the more usual Wigner 3j symbols, this is

$$\frac{1}{\sqrt{4\pi}} \left(\prod_{i=1}^3 c_{\ell_i}^i \sqrt{2\ell_i + 1} \right) \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ m_1 & m_2 & m_3 \end{pmatrix} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ -s_1 & -s_2 & -s_3 \end{pmatrix}.$$

One obtains the gradient (G) and curl (C) responses and cross responses summing the spin-weight responses ${}_{ab}\mathcal{R}$ in the following way ($a, b = \pm 1$):

$$\begin{aligned} \mathcal{R}^{GG} &= \frac{1}{2} ab {}_{ab}\mathcal{R}, & \mathcal{R}^{CC} &= \frac{1}{2} ab \mathcal{R} \\ \mathcal{R}^{GC} &= i \frac{a}{2} ab \mathcal{R}, & \mathcal{R}^{CG} &= -i \frac{b}{2} ab \mathcal{R}. \end{aligned} \quad (\text{A9})$$

In order to make progress, it is natural to split the F matrices into a diagonal and isotropic (m -independent), reference part F_{ℓ}^{iso} , and a perturbation δF . There are then two sorts of terms contributing to the change in response, proportional to $F^{\text{iso}} \delta F$ (and $\delta F F^{\text{iso}}$), and $\delta F \delta F$, where for small enough perturbations the former dominates. In that regime, the case of the cross-spectrum [Eq. (A5)] is particularly simple, and a fairly general result can be given, which we now describe.

From the addition rules of the Gaunt integrals, whenever two of the three legs of the pairs of matrices in Eq. (A8) are contracted (magnetic numbers summed over for the same multipole moments), the third leg numbers must then also match, and the entire sum results in a simple one-dimensional integral of Wigner small- d correlation functions (for example, the isotropic case contracts on the two CMB legs, resulting in the well-known isotropic lensing

response $\mathcal{R}_{LL'}^{MM'} \propto \delta_{MM'} \delta_{LL'} \mathcal{R}_L$). The $F^{\text{iso}} F$ terms in the case of the cross-spectrum [Eq. (A5)] also have two contracting legs (one CMB and the lensing leg), so that these terms are just as simple to evaluate, the difference to an isotropic response calculation being the replacement on the non-contracting CMB leg of F_ℓ^{iso} by the diagonal of the perturbation,

$$F_\ell^{\text{iso}} \rightarrow \frac{1}{(2\ell+1)} \sum_m \delta F_{\ell\ell}^{mm} (F^{\text{iso}} \delta F \text{ terms}). \quad (\text{A10})$$

For example, when just masking very small holes (in general positions), the perturbation to the F matrix in the absence of optimal filtering [as in Eq. (A4), but allowing for a nonazimuthally symmetric mask, and identifying the sky and fiducial transfer functions] may be written as

$$\delta F_{\ell\ell'}^{mm'} = -F_\ell^{\text{iso}} \frac{b'_\ell}{b_\ell} \sum_{\text{sources } i} \sigma_i Y_{\ell m}^\dagger(\hat{n}_i) Y_{\ell' m'}(\hat{n}_i), \quad (\text{A11})$$

where σ_i is the excised area for each small source. Using the rule in Eq. (A10), this results in the replacement

$$F_\ell^{\text{iso}} \rightarrow -F_\ell^{\text{iso}} \sum_{\text{sources } i} \left(\frac{\sigma_i}{4\pi} \right) \quad (\text{no optimal filtering}). \quad (\text{A12})$$

The prefactor is just the total excised sky fraction. The total cross-spectrum to the input lensing map is then reduced exactly by twice the excised sky fraction (since there is both an $F^{\text{iso}} \delta F$ and an $\delta F F^{\text{iso}}$ term). On the units of Fig. 17, this is identically -2.

With optimal filtering, the CMB is reconstructed within the holes, and the bias is much reduced. For the same population of sources, the matrix inverse in the optimal filtering F will introduce coupling between the sources contributions, but this can be neglected in the limit of small¹⁷ (not necessarily well separated) holes. The corresponding result becomes

$$F_\ell^{\text{iso}} \rightarrow -F_\ell^{\text{iso}} \left(1 - \frac{C_\ell^{\text{fid}}}{C_\ell^{\text{fid}} + N_r/b_\ell^2} \right) \sum_{\text{sources } i} \left(\frac{\sigma_i}{4\pi} \right) \quad (\text{optimal filtering}). \quad (\text{A13})$$

As expected, the bias to the cross-spectrum depends on the reconstruction signal to noise via the Wiener-filtering factor of the CMB reconstruction on the scales where the QE takes its most relevant contribution.

For larger holes, we cannot neglect the term quadratic in δF . One way to proceed is through a singular value

¹⁷This can be made more precise with the requirement that the maximal resolved multipole and hole radius are related through $\theta_{\text{mask}} \cdot \ell_{\text{max, noise-free}} \ll 2$.

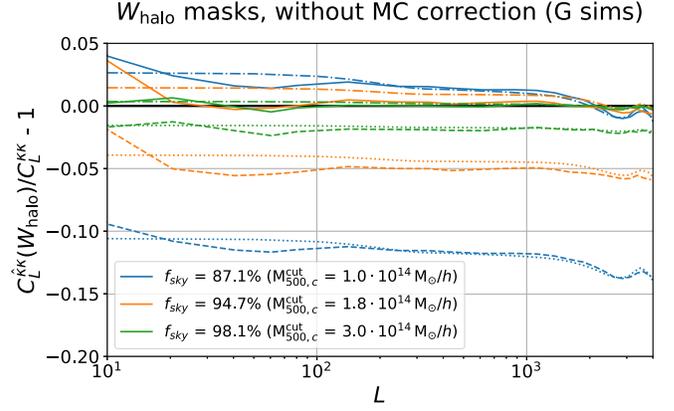


FIG. 18. Relative differences between the reconstructed cross $\hat{\kappa}\kappa$ -spectrum and the true WEBSKY $C_L^{\kappa\kappa}$ for three different uncorrelated W_{halo} masks normalizing with f_{sky}^{-1} to account for the mask area but no multiplicative Monte Carlo correction. Dashed lines show the same relative differences without the f_{sky}^{-1} normalization. Dot-dashed and dotted lines show the theoretical expectations for the model of Eq. (A5) with and without the f_{sky}^{-1} normalization, respectively. In the theoretical calculation, we assumed that all the disks have the same radius equal to the mean of the halo sizes (in terms of $2\theta_{500,c}$) in each mass-limited sample. For small holes, the f_{sky}^{-1} normalization overcorrects for the loss of power as discussed in Appendix A.

decomposition (SVD) of δF . Instead of contracting, the relevant leg separates in the large sum, and this also results in simpler one-dimensional integrals. This remains more complex, as one needs to calculate one set of integrals per singular value, per L and M . Hence, this approach is only practical when the perturbation to the filtering matrix can be described by a reasonable number of singular value. For example, a cluster mask of $5'$ radius only has ~ 10 relevant modes and is still perfectly doable. On the other hand, a Galactic mask like that of the *Planck*-lensing analysis has many thousands of these, and we have not attempted this calculation.

3. Comparison with Gaussian simulations

To validate the pipeline described in Sec. IV and to estimate the error bars in Figs. 8, 9, 11, and 12, we computed both the estimators in Eqs. (4.2) and (4.4) without f_L^{MC} correction, averaging over 20 G simulations that are uncorrelated to the fixed WEBSKY mask. As expected, W_X and W_X^{tot} masks give consistent results in the absence of mask correlations.

We also used the uncorrelated-mask simulation to compare with the analytic models for populations of small masked holes. We compare the relative differences between the reconstructed $\hat{\kappa}\kappa$ -spectrum and the true WEBSKY $\kappa\kappa$ -spectrum that was used to simulate the uncorrelated lensing fields, for three different W_{halo} masks, without the multiplicative Monte Carlo correction

defined in Eq. (4.5). Figure 18 shows that, as predicted by the analytic model, a f_{sky}^{-1} normalization would be a substantial overcorrection for masks made of many tiny holes that are partially refilled by the optimal filtering. Similar results have been found for all the other foreground masks.

APPENDIX B: EFFECTS OF MASKING NON-GAUSSIAN FIELDS

In this appendix, we explore the disagreement between the theory predictions based on simple Gaussian models and the simulation results for the threshold masks, W_{κ} , W_{CIB} , and W_y , as shown in Fig. 4.

To test higher-order non-Gaussian effects, as in Paper I, we constructed another set of foreground masks from the y , CIB, and κ fields by inverting their sign prior to the thresholding operation, $W_{-\text{CIB}}$, W_{-y} , and $W_{-\kappa}$. This is equivalent to masking the minima instead of the peaks of the foreground fields. This helps isolate higher-order effects that involve a higher even power of the perturbations compared to the expansion considered in our analytical model (where X and $-X$ masks give identical results). The results are shown in Fig. 19. We found good agreement only for $W_{-\kappa}$, while we recover at least a good qualitative trend for $W_{-\text{CIB}}$. However, for W_{-y} , the predictions instead lie away from the analytic curves, consistent with strongly non-Gaussian higher-order

effects (clearly visible in the mask shape difference shown in Fig. 20) not being captured in our Gaussian model. Here, the semiempirical results suggest only a very small bias when masking W_{-y} .

The theory predictions are based on full-sky spectra of non-Gaussian foreground maps and hence are sensitive to the peaks of the field inside the sky areas that get masked. Such features are enhanced for non-Gaussian fields compared to a purely Gaussian field. To test the impact of these non-Gaussian features, and to understand if they are responsible for the disagreement with the simulation results, we tried creating a Gaussianized version of the non-Gaussian foreground fields. To do this, we remapped the values in each pixel to the value that they would have in same the corresponding percentile in a Gaussian distribution having the same variance. Ideally, the masked foreground spectra after this procedure should be similar to the one obtained from non-Gaussian maps, but the full-sky ones used in Eq. (2.6) may be rather different, as now all the very bright extreme pixels have had their values scaled down. In Fig. 21, we show the analytic prediction based on these Gaussianized maps. We still do not find perfect agreement between simulations and theory predictions, but at least we are now able to reproduce the trends seen in the simulations: in particular, the W_y mask predictions now recover, at intermediate scales, distinctly different curves for masks retaining different f_{sky} s. Moreover, as expected, the semianalytic theory predictions and the pure Gaussian

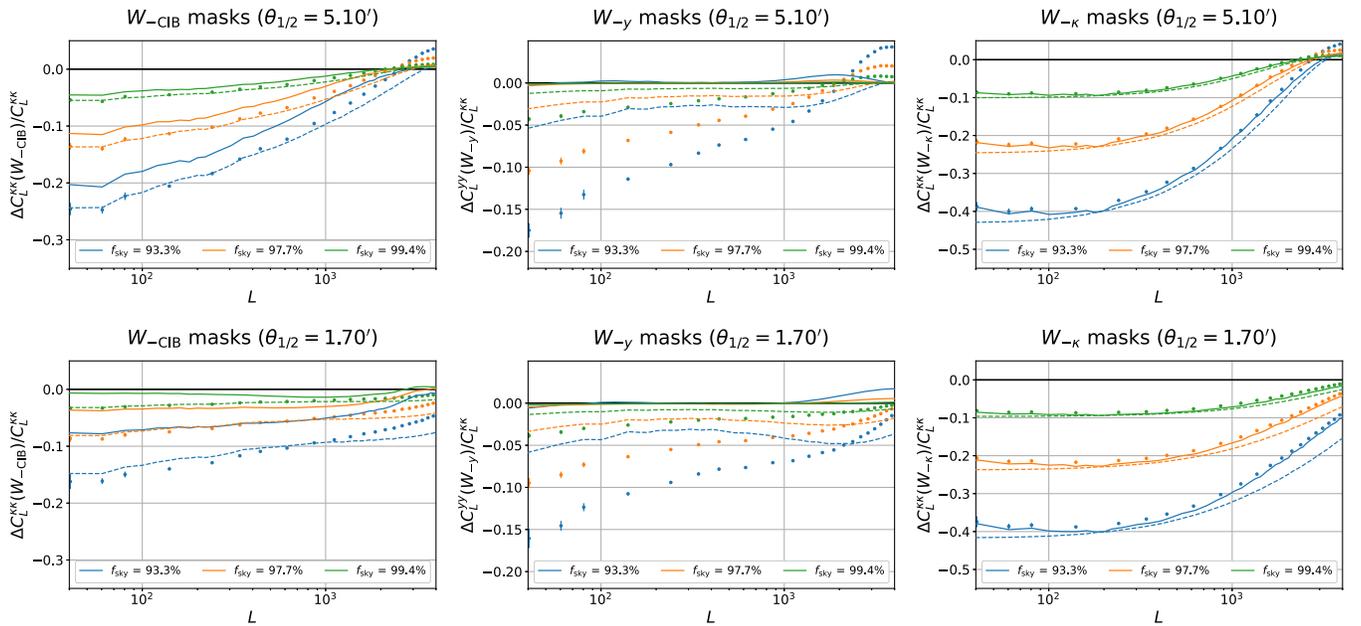


FIG. 19. Same as Fig. 4 for the sign-flipped masks $W_{-\text{CIB}}$, W_{-y} , and $W_{-\kappa}$ that remove minima of the foreground fields. Simulation measurements are shown as data points, the semianalytic Gaussian theory predictions are shown in solid, and the fully analytic Gaussian model is shown in dashed [evaluating the expectation values in Eq. (2.6) empirically from our simulations and analytically, respectively]. The top and bottom rows show the masks obtained by smoothing the foreground maps with a Gaussian beam of $\theta_{1/2} = 5.1'$, $1.7'$, respectively.

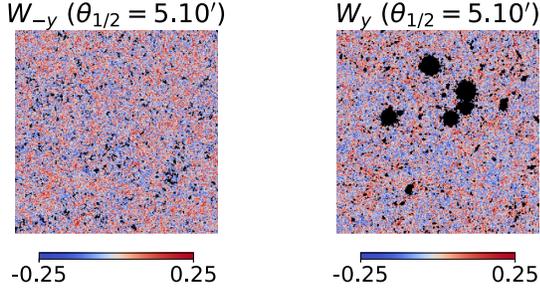


FIG. 20. WEBSKY κ field masked with W_y and W_{-y} masks obtained smoothing the foreground field with a $\theta_{1/2} = 5.1'$ Gaussian beam prior to thresholding. The masked fraction of the sky is the same for both masks ($f_{\text{sky}} = 93.3\%$). The tSZ signal is highly non-Gaussian and highly skewed, so that the mask covers large clusters in the standard case and a much larger number of small underdensities in the $-y$ case (for Gaussian fields, the two masks would be statistically identical).

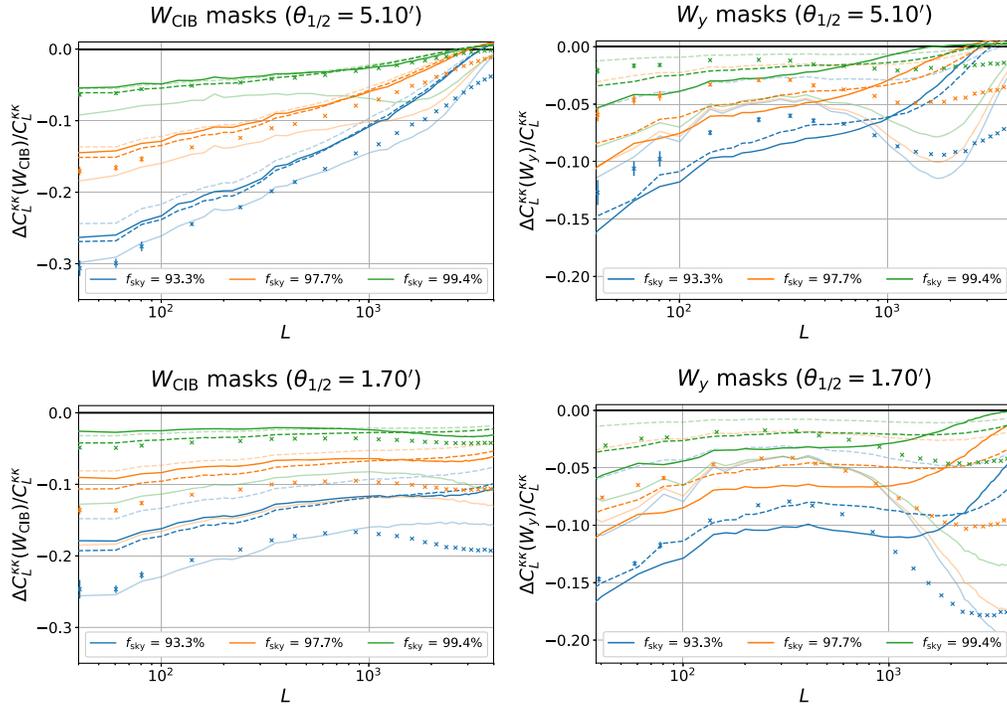


FIG. 21. Effect of LSS-correlated masking on the CMB lensing convergence power spectrum $C_L^{\kappa\kappa}$ as fractional difference between the $C_L^{\kappa\kappa}$ computed on the masked sky and its value computed on the full sky. We compare a subset of the simulation results shown in Fig. 4 with alternative analytical predictions. The semianalytical and analytical predictions of the Gaussian model obtained using the full-sky power spectra of Gaussianized foreground maps are shown in solid and in dashed, respectively. The baseline theory predictions (both semianalytical and fully analytical) based on non-Gaussian foreground maps and shown in Fig. 4 appear as faded lines. Top and bottom panels show results for the W_{CIB} and W_y masks constructed on foregrounds field smoothed with different Gaussian beams.

model obtained using the full-sky power spectra from Gaussianized maps are closer together (see dashed and solid lines).

Finally, to test the accuracy of the analytical model and to convince ourselves that discrepancies we observe between the analytical models and the simulations are simply due to non-Gaussianity of the foreground fields, we built correlated Gaussian map of the foreground fields, y^{G} , CIB^{G} , and κ^{G} , that preserve the cross-correlation between y , CIB, and κ as measured from the full-sky WEBSKY maps. We used these Gaussian realizations of the foreground fields to construct the Gaussian foreground W_y^{G} , $W_{\text{CIB}}^{\text{G}}$, and W_{κ}^{G} masks that we later applied on the corresponding Gaussian κ^{G} map. The results are shown in Fig. 22 and are in excellent agreement with the analytic model described in Sec. II B, confirming that the differences we see in the more realistic simulations are due to the non-Gaussian statistics.

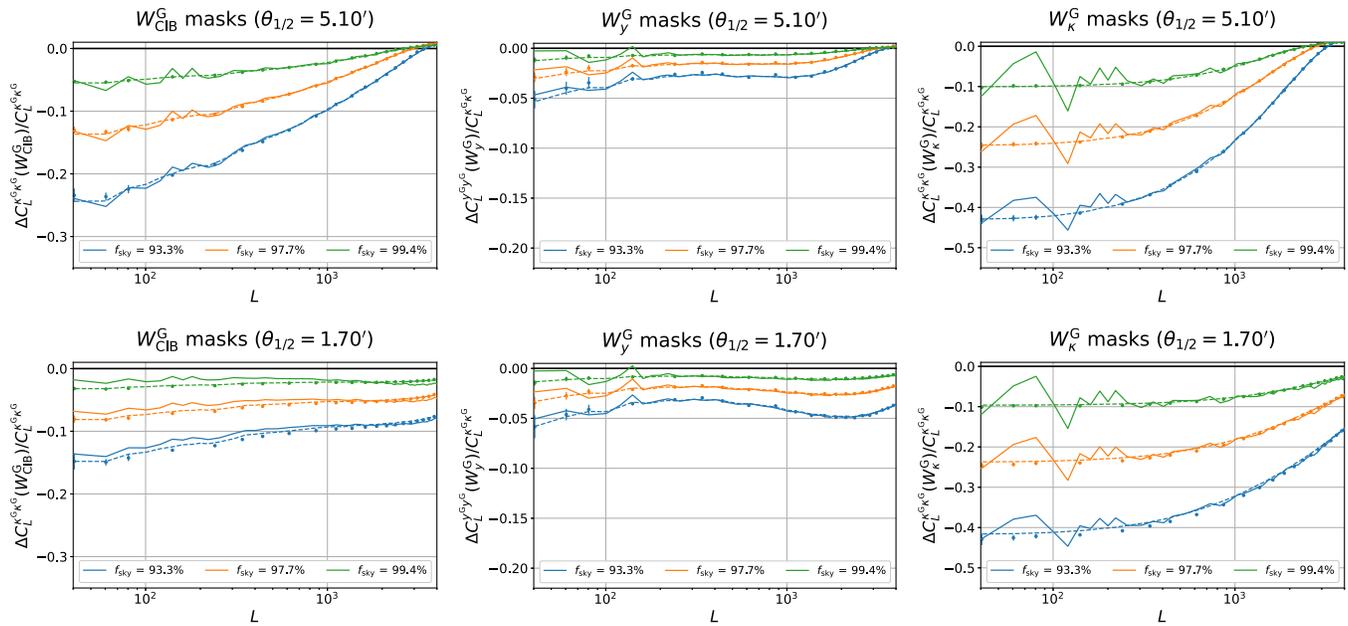


FIG. 22. Same as Fig. 19 for the W_y^G , W_{CIB}^G , and W_κ^G masks derived thresholding a Gaussian realization of the foreground fields and of CMB lensing keeping the same correlation between the components observed in the WEBSKY simulations.

- [1] P. A. R. Ade *et al.* (Planck Collaboration), Planck 2015 results. XXVII. The second Planck catalogue of Sunyaev-Zeldovich sources, *Astron. Astrophys.* **594**, A27 (2016).
- [2] L. E. Bleem *et al.* (SPT, DES Collaborations), The SPTpol extended cluster survey, *Astrophys. J. Suppl. Ser.* **247**, 25 (2020).
- [3] M. Hilton *et al.* (ACT, DES Collaborations), The Atacama Cosmology Telescope: A catalog of > 4000 Sunyaev-Zel'dovich galaxy clusters, *Astrophys. J. Suppl. Ser.* **253**, 3 (2021).
- [4] Gilbert P. Holder, Radio point sources and the thermal sz power spectrum, *Astrophys. J.* **580**, 36 (2002).
- [5] Masato Shirasaki, Impact of radio sources and cosmic infrared background on thermal Sunyaev-Zel'dovich—gravitational lensing cross-correlation, *Mon. Not. R. Astron. Soc.* **483**, 342 (2019).
- [6] Rupert Allison *et al.* (ACT Collaboration), The Atacama Cosmology Telescope: Measuring radio galaxy bias through cross-correlation with lensing, *Mon. Not. R. Astron. Soc.* **451**, 849 (2015).
- [7] Eli Dwek and Michael K. Barker, The cosmic radio and infrared backgrounds connection, *Astrophys. J.* **575**, 7 (2002).
- [8] P. A. R. Ade *et al.* (Planck Collaboration), Planck 2013 results. XVIII. The gravitational lensing-infrared background correlation, *Astron. Astrophys.* **571**, A18 (2014).
- [9] P. A. R. Ade *et al.* (Planck Collaboration), Planck 2015 results. XXIII. The thermal Sunyaev-Zeldovich effect—cosmic infrared background correlation, *Astron. Astrophys.* **594**, A23 (2016).
- [10] N. Aghanim *et al.* (Planck Collaboration), Planck 2015 results. XXII. A map of the thermal Sunyaev-Zeldovich effect, *Astron. Astrophys.* **594**, A22 (2016).
- [11] Yong-Seon Song, Asantha Cooray, Lloyd Knox, and Matias Zaldarriaga, The far-infrared background correlation with cmb lensing, *Astrophys. J.* **590**, 664 (2003).
- [12] Stephen J. Osborne, Duncan Hanson, and Olivier Doré, Extragalactic foreground contamination in temperature-based CMB lens reconstruction, *J. Cosmol. Astropart. Phys.* **03** (2014) 024.
- [13] A. van Engelen, S. Bhattacharya, N. Sehgal, G. P. Holder, O. Zahn, and D. Nagai, CMB lensing power spectrum biases from galaxies and clusters using high-angular resolution temperature maps, *Astrophys. J.* **786**, 13 (2014).
- [14] Toshiya Namikawa, Duncan Hanson, and Ryuichi Takahashi, Bias-hardened CMB lensing, *Mon. Not. R. Astron. Soc.* **431**, 609 (2013).
- [15] Emmanuel Schaan and Simone Ferraro, Foreground-Immune CMB Lensing with Shear-Only Reconstruction, *Phys. Rev. Lett.* **122**, 181301 (2019).
- [16] Omar Darwish *et al.*, The Atacama Cosmology Telescope: A CMB lensing mass map over 2100 square degrees of sky

- and its cross-correlation with BOSS-CMASS galaxies, *Mon. Not. R. Astron. Soc.* **500**, 2250 (2020).
- [17] Mathew S. Madhavacheril and J. Colin Hill, Mitigating foreground biases in CMB lensing reconstruction using cleaned gradients, *Phys. Rev. D* **98**, 023534 (2018).
- [18] Giulio Fabbian, Antony Lewis, and Dominic Beck, CMB lensing reconstruction biases in cross-correlation with large-scale structure probes, *J. Cosmol. Astropart. Phys.* **10** (2019) 057.
- [19] Sigurd Naess *et al.*, The Atacama Cosmology Telescope: Arcminute-resolution maps of 18 000 square degrees of the microwave sky from ACT 2008–2018 data combined with Planck, *J. Cosmol. Astropart. Phys.* **12** (2020) 046.
- [20] G. Fabbian, J. Carron, A. Lewis, and M. Lembo, Lensed CMB power spectrum biases from masking extragalactic sources, *Phys. Rev. D* **103**, 043535 (2021).
- [21] N. Aghanim *et al.* (Planck Collaboration), Planck 2018 results. VIII. Gravitational lensing, *Astron. Astrophys.* **641**, A8 (2020).
- [22] Wayne Hu and Takemi Okamoto, Mass reconstruction with CMB polarization, *Astrophys. J.* **574**, 566 (2002).
- [23] Michael Kesden, Asantha Cooray, and Marc Kamionkowski, Lensing reconstruction with CMB temperature and polarization, *Phys. Rev. D* **67**, 123507 (2003).
- [24] Vanessa Böhm, Marcel Schmittfull, and Blake D. Sherwin, Bias to CMB lensing measurements from the bispectrum of large-scale structure, *Phys. Rev. D* **94**, 043519 (2016).
- [25] Geraint Pratten and Antony Lewis, Impact of post-Born lensing on the CMB, *J. Cosmol. Astropart. Phys.* **08** (2016) 047.
- [26] Vanessa Böhm, Blake D. Sherwin, Jia Liu, J. Colin Hill, Marcel Schmittfull, and Toshiya Namikawa, Effect of non-Gaussian lensing deflections on CMB lensing measurements, *Phys. Rev. D* **98**, 123510 (2018).
- [27] Dominic Beck, Giulio Fabbian, and Josquin Errard, Lensing reconstruction in post-Born cosmic microwave background weak lensing, *Phys. Rev. D* **98**, 043512 (2018).
- [28] Jia Liu, J. Colin Hill, Blake D. Sherwin, Andrea Petri, Vanessa Böhm, and Zoltán Haiman, CMB lensing beyond the power spectrum: Cosmological constraints from the one-point probability distribution function and peak counts, *Phys. Rev. D* **94**, 103501 (2016).
- [29] Alberto Vallinotto, Using CMB lensing to constrain the multiplicative bias of cosmic shear, *Astrophys. J.* **759**, 32 (2012).
- [30] Emmanuel Schaan, Elisabeth Krause, Tim Eifler, Olivier Doré, Hironao Miyatake, Jason Rhodes, and David N. Spergel, Looking through the same lens: Shear calibration for LSST, Euclid, and WFIRST with stage 4 CMB lensing, *Phys. Rev. D* **95**, 123512 (2017).
- [31] Ross Cawthon, Effects of redshift uncertainty on cross-correlations of CMB lensing and galaxy surveys, *Phys. Rev. D* **101**, 063509 (2020).
- [32] Emmanuel Schaan, Simone Ferraro, and Uroš Seljak, Photo-z outlier self-calibration in weak lensing surveys, *J. Cosmol. Astropart. Phys.* **12** (2020) 001.
- [33] S. Ilić *et al.* (Euclid Collaboration), *Euclid* preparation: XV. Forecasting cosmological constraints for the *Euclid* and CMB joint analysis, *Astron. Astrophys.* **657**, A91 (2022).
- [34] Noah Sailer, Emanuele Castorina, Simone Ferraro, and Martin White, Cosmology at high redshift—a probe of fundamental physics, *J. Cosmol. Astropart. Phys.* **12** (2021) 049.
- [35] J. Colin Hill and David N. Spergel, Detection of thermal SZ-CMB lensing cross-correlation in Planck nominal mission data, *J. Cosmol. Astropart. Phys.* **02** (2014) 030.
- [36] N. Battaglia, J. C. Hill, and N. Murray, Deconstructing thermal Sunyaev–Zel’dovich—gravitational lensing cross-correlations: Implications for the intracluster medium, *Astrophys. J.* **812**, 154 (2015).
- [37] Alireza Hojjati, Ian G. McCarthy, Joachim Harnois-Deraps, Yin-Zhe Ma, Ludovic Van Waerbeke, Gary Hinshaw, and Amandine M. C. Le Brun, Dissecting the thermal Sunyaev-Zeldovich-gravitational lensing cross-correlation with hydrodynamical simulations, *J. Cosmol. Astropart. Phys.* **10** (2015) 047.
- [38] Ludovic Van Waerbeke, Gary Hinshaw, and Norman Murray, Detection of warm and diffuse baryons in large scale structure from the cross-correlation of gravitational lensing and the thermal Sunyaev-Zeldovich effect, *Phys. Rev. D* **89**, 023508 (2014).
- [39] Ken Osato, Samuel Flender, Daisuke Nagai, Masato Shirasaki, and Naoki Yoshida, Investigating cluster astrophysics and cosmology with cross-correlation of the thermal Sunyaev–Zel’dovich effect and weak lensing, *Mon. Not. R. Astron. Soc.* **475**, 532 (2018).
- [40] Matthieu Bethermin, Olivier Dore, and Guilaine Lagache, Where stars form and live at high redshift: Clues from the infrared, *Astron. Astrophys.* **537**, L5 (2012).
- [41] A. Maniyar, M. Béthermin, and G. Lagache, Simple halo model formalism for the cosmic infrared background and its correlation with the thermal Sunyaev Zel’dovich effect, *Astron. Astrophys.* **645**, A40 (2021).
- [42] P. A. R. Ade *et al.* (Planck Collaboration), Planck 2013 results. XVIII. The gravitational lensing-infrared background correlation, *Astron. Astrophys.* **571**, A18 (2014).
- [43] P. Ade *et al.* (Simons Observatory Collaboration), The Simons observatory: Science goals and forecasts, *J. Cosmol. Astropart. Phys.* **02** (2019) 056.
- [44] Manuel Aravena *et al.*, The CCAT-prime submillimeter observatory, [arXiv:1909.02587](https://arxiv.org/abs/1909.02587).
- [45] Fiona McCarthy and Mathew S. Madhavacheril, Improving models of the cosmic infrared background using CMB lensing mass maps, *Phys. Rev. D* **103**, 103515 (2021).
- [46] Blake D. Sherwin and Marcel Schmittfull, Delensing the CMB with the cosmic infrared background, *Phys. Rev. D* **92**, 043005 (2015).
- [47] Byeonghee Yu, J. Colin Hill, and Blake D. Sherwin, Multitracer CMB delensing maps from Planck and WISE data, *Phys. Rev. D* **96**, 123511 (2017).
- [48] Dominic Beck, Josquin Errard, and Radek Stompor, Impact of polarized galactic foreground emission on CMB lensing reconstruction and delensing of B-modes, *J. Cosmol. Astropart. Phys.* **06** (2020) 030.
- [49] Kevork Abazajian *et al.*, CMB-S4 science case, reference design, and project plan, [arXiv:1907.04473](https://arxiv.org/abs/1907.04473).

- [50] Asantha Cooray and Ravi K. Sheth, Halo models of large scale structure, *Phys. Rep.* **372**, 1 (2002).
- [51] J. F. Navarro, C. S. Frenk, and S. D. M. White, A universal density profile from hierarchical clustering, *Astrophys. J.* **490**, 493 (1997).
- [52] Alan R. Duffy, Joop Schaye, Scott T. Kay, and Claudio Dalla Vecchia, Dark matter halo concentrations in the Wilkinson Microwave Anisotropy Probe year 5 cosmology, *Mon. Not. R. Astron. Soc.* **390**, L64 (2008).
- [53] Ravi K. Sheth and Giuseppe Tormen, Large-scale bias and the peak background split, *Mon. Not. R. Astron. Soc.* **308**, 119 (1999).
- [54] Aditya Rotti, Boris Bolliet, Jens Chluba, and Mathieu Remazeilles, Removing the giants and learning from the crowd: A new SZ power spectrum method and revised Compton y -map analysis, *Mon. Not. R. Astron. Soc.* **503**, 5310 (2021).
- [55] Kendrick M. Smith, Mathew S. Madhavacheril, Moritz Münchmeyer, Simone Ferraro, Utkarsh Giri, and Matthew C. Johnson, KSZ tomography and the bispectrum, [arXiv:1810.13423](https://arxiv.org/abs/1810.13423).
- [56] Ravi K. Sheth, H. J. Mo, and Giuseppe Tormen, Ellipsoidal collapse and an improved model for the number and spatial distribution of dark matter haloes, *Mon. Not. R. Astron. Soc.* **323**, 1 (2001).
- [57] Ravi K. Sheth and Giuseppe Tormen, An excursion set model of hierarchical clustering: Ellipsoidal collapse and the moving barrier, *Mon. Not. R. Astron. Soc.* **329**, 61 (2002).
- [58] Jeremy L. Tinker, Brant E. Robertson, Andrey V. Kravtsov, Anatoly Klypin, Michael S. Warren, Gustavo Yepes, and Stefan Gottlöber, The large-scale bias of dark matter halos: Numerical calibration and model tests, *Astrophys. J.* **724**, 878 (2010).
- [59] George Stein, Marcelo A. Alvarez, J. Richard Bond, Alexander van Engelen, and Nicholas Battaglia, The WEBSKY extragalactic CMB simulations, *J. Cosmol. Astropart. Phys.* **10** (2020) 012.
- [60] George Stein, Marcelo A. Alvarez, and J. Richard Bond, The mass-Peak Patch algorithm for fast generation of deep all-sky dark matter halo catalogues and its N-Body validation, *Mon. Not. R. Astron. Soc.* **483**, 2236 (2019).
- [61] Linda Blot *et al.*, Comparing approximate methods for mock catalogues and covariance matrices II: Power spectrum multipoles, *Mon. Not. R. Astron. Soc.* **485**, 2806 (2019).
- [62] Manuel Colavincenzo *et al.*, Comparing approximate methods for mock catalogues and covariance matrices—III: bispectrum, *Mon. Not. R. Astron. Soc.* **482**, 4883 (2019).
- [63] Martha Lippich *et al.*, Comparing approximate methods for mock catalogues and covariance matrices—I. Correlation function, *Mon. Not. R. Astron. Soc.* **482**, 1786 (2019).
- [64] Neelima Sehgal, Paul Bode, Sudeep Das, Carlos Hernandez-Monteagudo, Kevin Huffenberger, Yen-Ting Lin, Jeremiah P. Ostriker, and Hy Trac, Simulations of the microwave sky, *Astrophys. J.* **709**, 920 (2010).
- [65] Zack Li, Giuseppe Puglisi, Mathew S. Madhavacheril, and Marcelo A. Alvarez, Simulated catalogs and maps of radio galaxies at millimeter wavelengths in WEBSKY, [arXiv:2110.15357](https://arxiv.org/abs/2110.15357).
- [66] Vanessa Böhm, Chirag Modi, and Emanuele Castorina, Lensing corrections on galaxy-lensing cross correlations and galaxy-galaxy auto correlations, *J. Cosmol. Astropart. Phys.* **03** (2020) 045.
- [67] Giulio Fabbian, Matteo Calabrese, and Carmelita Carbone, CMB weak-lensing beyond the Born approximation: A numerical approach, *J. Cosmol. Astropart. Phys.* **02** (2018) 050.
- [68] E. Hivon, K. M. Gorski, C. B. Netterfield, B. P. Crill, S. Prunet, and F. Hansen, MASTER of the CMB anisotropy power spectrum: A fast method for statistical analysis of large and complex CMB data sets, *Astrophys. J.* **567**, 2 (2002).
- [69] David Alonso, Javier Sanchez, and Anže Slosar (LSST Dark Energy Science Collaboration), A unified pseudo- C_ℓ framework, *Mon. Not. R. Astron. Soc.* **484**, 4127 (2019).
- [70] Ryu Makiya, Shin'ichiro Ando, and Eiichiro Komatsu, Joint analysis of the thermal Sunyaev–Zeldovich effect and 2MASS galaxies: Probing gas physics in the local Universe and beyond, *Mon. Not. R. Astron. Soc.* **480**, 3928 (2018).
- [71] N. Gupta *et al.* (SPT Collaboration), Fractional polarization of extragalactic sources in the 500 deg² SPTpol survey, *Mon. Not. R. Astron. Soc.* **490**, 5712 (2019).
- [72] N. Aghanim *et al.* (Planck Collaboration), Planck 2018 results. VIII. Gravitational lensing, *Astron. Astrophys.* **641**, A8 (2020).
- [73] Takemi Okamoto and Wayne Hu, CMB lensing reconstruction on the full sky, *Phys. Rev. D* **67**, 083002 (2003).
- [74] Srinivasan Raghunathan, Nathan Whitehorn, Marcelo A. Alvarez, Han Aung, Nicholas Battaglia, Gilbert P. Holder, Daisuke Nagai, Elena Pierpaoli, Christian L. Reichardt, and Joaquin D. Vieira, Constraining cluster virialization mechanism and cosmology using thermal-SZ-selected clusters from future CMB surveys, *Astrophys. J.* **926**, 172 (2022).
- [75] Kendrick M. Smith, Oliver Zahn, and Olivier Dore, Detection of gravitational lensing in the cosmic microwave background, *Phys. Rev. D* **76**, 043510 (2007).
- [76] Mark Mirmelstein, Julien Carron, and Antony Lewis, Optimal filtering for CMB lensing reconstruction, *Phys. Rev. D* **100**, 123509 (2019).
- [77] Duncan Hanson, Anthony Challinor, George Efstathiou, and Pawel Bielewicz, CMB temperature lensing power reconstruction, *Phys. Rev. D* **83**, 043005 (2011).
- [78] Antony Lewis, Anthony Challinor, and Duncan Hanson, The shape of the CMB lensing bispectrum, *J. Cosmol. Astropart. Phys.* **03** (2011) 018.
- [79] Toshiya Namikawa, Duncan Hanson, and Ryuichi Takahashi, Bias-hardened CMB lensing, *Mon. Not. R. Astron. Soc.* **431**, 609 (2013).
- [80] S. Adachi *et al.* (POLARBEAR Collaboration), Internal Delensing of Cosmic Microwave Background Polarization B -Modes with the POLARBEAR Experiment, *Phys. Rev. Lett.* **124**, 131301 (2020).
- [81] Dongwon Han *et al.* (ACT Collaboration), The Atacama Cosmology Telescope: Delensed power spectra and parameters, *J. Cosmol. Astropart. Phys.* **01** (2021) 031.
- [82] Giulio Fabbian and Radek Stompor, High precision simulations of weak lensing effect on Cosmic Microwave Background polarization, *Astron. Astrophys.* **556**, A109 (2013).

- [83] A. Manzotti *et al.* (Herschel, SPT Collaborations), CMB polarization B-mode delensing with SPTpol and Herschel, *Astrophys. J.* **846**, 45 (2017).
- [84] Patricia Larsen, Anthony Challinor, Blake D. Sherwin, and Daisy Mak, Demonstration of Cosmic Microwave Background Delensing Using the Cosmic Infrared Background, *Phys. Rev. Lett.* **117**, 151102 (2016).
- [85] P. A. R. Ade *et al.* (Planck Collaboration), Planck 2015 results. XXIII. The thermal Sunyaev-Zeldovich effect–cosmic infrared background correlation, *Astron. Astrophys.* **594**, A23 (2016).
- [86] Mathew S. Madhavacheril, Nicholas Battaglia, and Hironao Miyatake, Fundamental physics from future weak-lensing calibrated Sunyaev-Zeldovich galaxy cluster counts, *Phys. Rev. D* **96**, 103525 (2017).
- [87] Srinivasan Raghunathan, Sanjaykumar Patil, Eric J. Baxter, Federico Bianchini, Lindsey E. Bleem, Thomas M. Crawford, Gilbert P. Holder, Alessandro Manzotti, and Christian L. Reichardt, Measuring galaxy cluster masses with CMB lensing using a maximum likelihood estimator: Statistical and systematic error budgets for future experiments, *J. Cosmol. Astropart. Phys.* **08** (2017) 030.
- [88] Hironao Miyatake *et al.*, Weak-lensing mass calibration of ACTPol Sunyaev-Zeldovich clusters with the hyper supprime-cam survey, *Astrophys. J.* **875**, 63 (2019).
- [89] Srinivasan Raghunathan, Federico Bianchini, and Christian L. Reichardt, Imprints of gravitational lensing in the Planck cosmic microwave background data at the location of WISExSCOS galaxies, *Phys. Rev. D* **98**, 043506 (2018).
- [90] S. Raghunathan *et al.* (DES Collaboration), Mass calibration of optically selected DES clusters using a measurement of CMB-cluster lensing with SPTpol data, *Astrophys. J.* **872**, 170 (2019).
- [91] Uros Seljak and Matias Zaldarriaga, Lensing induced cluster signatures in cosmic microwave background, *Astrophys. J.* **538**, 57 (2000).
- [92] Scott Dodelson, CMB-cluster lensing, *Phys. Rev. D* **70**, 023009 (2004).
- [93] Gilbert P. Holder and Arthur Kosowsky, Gravitational lensing of the microwave background by galaxy clusters, *Astrophys. J.* **616**, 8 (2004).
- [94] Antony Lewis and Lindsay King, Cluster masses from CMB and galaxy weak lensing, *Phys. Rev. D* **73**, 063006 (2006).
- [95] Wayne Hu, Simon DeDeo, and Chris Vale, Cluster mass estimators from cmb temperature and polarization lensing, *New J. Phys.* **9**, 441 (2007).
- [96] Srinivasan Raghunathan, Gilbert P. Holder, James G. Bartlett, Sanjaykumar Patil, Christian L. Reichardt, and Nathan Whitehorn, An inpainting approach to tackle the kinematic and thermal SZ induced biases in CMB-cluster lensing estimators, *J. Cosmol. Astropart. Phys.* **11** (2019) 037.
- [97] Matteo Maturi, Matthias Bartelmann, Massimo Meneghetti, and Lauro Moscardini, Gravitational lensing of the CMB by galaxy clusters, *Astron. Astrophys.* **436**, 37 (2005).
- [98] Jaiyul Yoo and Matias Zaldarriaga, Improved estimation of cluster mass profiles from the cosmic microwave background, *Phys. Rev. D* **78**, 083002 (2008).
- [99] P. A. R. Ade *et al.* (Planck Collaboration), Planck 2015 results. XXVII. The second Planck catalogue of Sunyaev-Zeldovich sources, *Astron. Astrophys.* **594**, A27 (2016).
- [100] Sanjaykumar Patil, Srinivasan Raghunathan, and Christian L. Reichardt, Suppressing the thermal SZ-induced variance in CMB-cluster lensing estimators, *Astrophys. J.* **888**, 9 (2020).
- [101] C. Reuter *et al.*, The complete redshift distribution of dusty star-forming galaxies from the SPT-SZ survey, *Astrophys. J.* **902**, 78 (2020).
- [102] Zhen-Yi Cai, Andrea Lapi, Jun-Qing Xia, Gianfranco De Zotti, Mattia Negrello, Carlotta Gruppioni, Emma Rigby, Guillaume Castex, Jacques Delabrouille, and Luigi Danese, A hybrid model for the evolution of galaxies and Active Galactic Nuclei in the infrared, *Astrophys. J.* **768**, 21 (2013).
- [103] A. Lapi, M. Negrello, J. Gonzalez-Nuevo, Z. Y. Cai, G. De Zotti, and L. Danese, Effective models for statistical studies of galaxy-scale gravitational lensing, *Astrophys. J.* **755**, 46 (2012).
- [104] Marcelo Alvarez, Federico Bianchini, Giulio Fabbian, Mattia Negrello, and Giuseppe Puglisi (private communication).
- [105] Steve K. Choi *et al.*, Sensitivity of the Prime-Cam instrument on the CCAT-prime telescope, *J. Low Temp. Phys.* **199**, 1089 (2020).
- [106] Christopher M. Hirata and Uros Seljak, Analyzing weak lensing of the cosmic microwave background using the likelihood function, *Phys. Rev. D* **67**, 043001 (2003).
- [107] Christopher M. Hirata and Uros Seljak, Reconstruction of lensing from the cosmic microwave background polarization, *Phys. Rev. D* **68**, 083002 (2003).
- [108] Julien Carron and Antony Lewis, Maximum *a posteriori* CMB lensing reconstruction, *Phys. Rev. D* **96**, 063510 (2017).
- [109] Marcelo Ponce, Ramses van Zon, Scott Northrup, Daniel Gruner, Joseph Chen, Fatih Ertinaz, Alexey Fedoseev, Leslie Groer, Fei Mao, Bruno C. Mundim, Mike Nolta, Jaime Pinto, Marco Saldarriaga, Vladimir Slavnic, Erik Spence, Ching-Hsing Yu, and W. Richard Peltier, Deploying a top-100 supercomputer for large parallel workloads: The Niagara supercomputer, [arXiv:1907.13600](https://arxiv.org/abs/1907.13600).
- [110] W. L. K. Wu *et al.*, A measurement of the cosmic microwave background lensing potential and power spectrum from 500 deg² of SPTpol temperature and polarization data, *Astrophys. J.* **884**, 70 (2019).
- [111] Julien Carron, Position-space curved-sky anisotropy quadratic estimation, [arXiv:1908.02016](https://arxiv.org/abs/1908.02016).
- [112] Benjamin D. Wandelt, Eric Hivon, and Krzysztof M. Gorski, The Pseudo- C_l method: Cosmic microwave background anisotropy power spectrum statistics for high precision cosmology, *Phys. Rev. D* **64**, 083003 (2001).
- [113] Anthony Challinor and Gayoung Chon, Geometry of weak lensing of CMB polarization, *Phys. Rev. D* **66**, 127301 (2002).