Mitigating foreground bias to the CMB lensing power spectrum for a CMB-HD survey

Dongwon Han and Neelima Sehgal

Physics and Astronomy Department, Stony Brook University, Stony Brook, New York 11794, USA

A promising way to measure the distribution of matter on small scales ($k \sim 10 \, h \, \mathrm{Mpc}^{-1}$) is to use gravitational lensing of the cosmic microwave background (CMB). CMB-HD, a proposed high-resolution, low-noise millimeter survey over half the sky, can measure the CMB lensing auto spectrum on such small scales enabling measurements that can distinguish between a cold dark matter (CDM) model and alternative models designed to solve problems with CDM on small scales. However, extragalactic foregrounds can bias the CMB lensing auto spectrum if left untreated. We present a foreground mitigation strategy that provides a path to reduce the bias from two of the most dominant foregrounds, the thermal Sunyaev-Zel'dovich effect and the cosmic infrared background. Given the level of realism included in our analysis, we find that the thermal Sunyaev-Zel'dovich alone and the cosmic infrared background alone bias the lensing auto spectrum by 0.6σ and 1.1σ, respectively, in the lensing multipole range of $L \in [5000, 20000]$ for a CMB-HD survey; combined these foregrounds yield a bias of only 1.3σ. Including these foregrounds, we also find that a CMB-HD survey can distinguish between a CDM model and a 10^{-22} eV fuzzy dark matter model at the 5σ level. These results provide an important step in demonstrating that foreground contamination can be sufficiently reduced to enable a robust measurement of the small-scale matter power spectrum with CMB-HD.

I. INTRODUCTION

The nature of dark matter (DM) is a major open question in cosmology and particle physics. In the context of cosmology, a cold dark matter (CDM) model, where dark matter consists of a nonrelativistic, collisionless particle, yields predictions that are consistent with observations of the large scale structure (e.g., [1–6]). However, on smaller scales less than 10 kpc (corresponding to $M < 10^9 \, M_\odot$), CDM may overestimate the amount of structure (e.g., [7–17]). Many alternative models of dark matter have been proposed to solve this issue. These alternative models include warm dark matter [18–20], fuzzy dark matter (FDM) [21–30], self-interacting dark matter [31–39], and superfluid dark matter [40,41], to name a few. These alternative DM models predict varying amounts of the suppression of structure growth on small scales. In addition, baryonic processes, such as Active Galactic Nuclei (AGN) feedback, can alter the distribution of matter on small scales (e.g., [13,42–44]). Recent work suggests that baryons alter the shape of the matter power spectrum in a finite set of ways that can be described with a small set of parameters [45,46]; moreover, alternate dark matter models appear to change the shape of the matter power spectrum in a way that differs from baryonic processes [47]. Therefore, a robust measurement of the matter power spectrum below 10 kpc scales will impose stringent constraints on

alternative dark matter models and baryonic effects, and has the potential to distinguish between the two.

One promising avenue to measure small-scale structure is a high-resolution (0.25 arcminute), low-noise (0.5 μK-arcmin) cosmic microwave background (CMB) lensing measurement [47–50]. Unlike baryonic tracers, such as a galaxy number counts or Lyman-α, CMB lensing measurements probe the underlying dark matter distribution directly through gravitational lensing. Furthermore, in a cold dark matter paradigm, lower mass halos formed first at high redshift. Since the CMB lensing kernel peaks at relatively high redshift ($z \approx 2$), it is, in principle, more sensitive to the suppression of structure than local probes. CMB-HD, an ultradeep, high-resolution millimeter-wave survey over half the sky, has the potential to measure the CMB lensing auto spectrum on small scales of order $k \sim 10 \, h \, \mathrm{Mpc}^{-1}$, which no precursor CMB survey can achieve [48–50]. In particular, CMB-HD is sensitive enough to make a robust measurement of the small-scale matter power spectrum and determine whether baryonic physics or modifications to CDM, or potentially both, are needed to match observations of the matter distribution on small scales.

The CMB lensing signal on small scales gains most of its signal-to-noise ratio (SNR) from CMB temperature maps [47]. However, CMB temperature maps include contamination from extragalactic foregrounds, such as the thermal

and kinetic Sunyaev-Zel'dovich effects (TSZ and KSZ), the cosmic infrared background (CIB), and radio galaxies (radio). If untreated, these extragalactic foreground components can significantly bias the CMB lensing signal [51–56]. Many techniques have been proposed to mitigate this foreground bias, such as gradient-cleaning, foreground deprojection, and biased-hardened lensing estimators [55,57,58]. However, previous works have mainly focused on the impact of foregrounds on larger scales ($L \leq 5000$). In this work, we focus on a mitigation strategy for smaller scales ($L \in [5000, 20000]$), which contribute most of the SNR in distinguishing, for example, alternative DM models from a CDM model [47]. We apply these techniques to simulations from [59] that include realistic non-Gaussian extragalactic foregrounds, after modifying these simulations as described in Sec. II. In Sec. III, we describe the CMB lensing estimator and foreground mitigation techniques used in this work. In Sec. IV, we show that we are able to reduce the two most dominant foregrounds, the TSZ and CIB, enough to make an unbiased measurement of the CMB lensing auto spectrum in the range of $L \in [5000, 20000]$, and robustly distinguish between interesting dark matter models. We discuss additional systematic considerations in Sec. V, and conclude in Sec. VI.

II. SIMULATIONS

We use the non-Gaussian extragalactic foreground simulations from [59] (hereafter S10 simulations), since these simulations have extragalactic foregrounds correlated with each other. We rescale the S10 CIB and TSZ maps by a factor of 0.75 following [51,60] to make the S10 simulations more closely match the latest observations [61–65]. We cut out four $20° \times 20° = 400$ square degree patches near the equator, and reproject the corresponding S10 CIB,

TSZ, and radio maps from this patch onto a flat-sky grid with 0.25 arcmin resolution to match the resolution of CMB-HD at 150 GHz. We call these maps the S10 CIB, TSZ, and radio patch maps.

A. Creation of high-resolution lensing potential map correlated with TSZ and CIB

The correlation between the CIB and the lensing convergence map, κ, in the S10 simulations is roughly 35% at 150 GHz [66], which is significantly lower than the 70% correlation measured subsequently by *Planck* [67]. In addition, the S10κ map, which has a resolution of 1 arcminute, only has the information up to $\ell \approx 10000$, missing the small scale information needed for this analysis. Therefore, we replace the S10κ map with one we make ourselves that has higher resolution and 70% correlation with the CIB by construction. Similarly, the TSZ-κ anticorrelation is 45% at 150 GHz in S10, as opposed to roughly 50% found in [68] over the range of $\ell \in [100, 10000]$. Thus we increase the correlation between the S10 TSZ map and our new κ map to 50%. In this work, we focus on the bias from the TSZ and CIB to the lensing signal, as these foregrounds generally have the largest contributions to the lensing bias [51,66]; thus we do not explicitly include KSZ-κ and radio-κ correlations in the construction of the new κ map. However, we do include the KSZ and radio signals as additional sources of noise throughout this work. We leave an in-depth exploration of KSZ and radio bias to the lensing signal for future work.

We model this new κ map as a sum of two parts: one that is correlated with the foregrounds and one that is independent of the foregrounds. We create this new κ map following Eq. (1) below.

$$\kappa = M_G + \frac{1}{1 - [\rho^{\mathrm{CIB \times TSZ}}(\ell)]^2} [\rho^{\kappa \times \mathrm{CIB}}(\ell) - \rho^{\mathrm{CIB \times TSZ}}(\ell)\rho^{\kappa \times \mathrm{TSZ}}(\ell)] \sqrt{\frac{C_\ell^{\kappa\kappa}}{C_\ell^{\mathrm{CIB}}}} M_{\mathrm{CIB}}^{S10}$$

$$+ \frac{1}{1 - [\rho^{\mathrm{CIB \times TSZ}}(\ell)]^2} [\rho^{\kappa \times \mathrm{TSZ}}(\ell) - \rho^{\mathrm{CIB \times TSZ}}(\ell)\rho^{\kappa \times \mathrm{CIB}}(\ell)] \sqrt{\frac{C_\ell^{\kappa\kappa}}{C_\ell^{\mathrm{TSZ}}}} M_{\mathrm{TSZ}}^{S10} \tag{1}$$

Here, $C_\ell^{\kappa\kappa}$ is a theory lensing convergence power spectrum generated using CAMB [69]. M_{CIB}^{S10} and M_{TSZ}^{S10} are the S10 CIB and TSZ patch maps discussed above. We calculate C_ℓ^{CIB} and C_ℓ^{TSZ} by measuring the power spectrum of the full-sky S10 CIB and TSZ maps at 150 GHz. These raw spectra are noisy due to the single realization of the S10 simulation set; thus, we smooth the resulting power spectra with a 20σ Gaussian kernel, which smooths the spectra over $\Delta\ell \sim 20$, to reduce this noise. Similarly, we compute the correlation coefficient, $\rho^{\mathrm{CIB \times TSZ}}(\ell) = C_\ell^{\mathrm{CIB \times TSZ}} / \sqrt{C_\ell^{\mathrm{CIB}} C_\ell^{\mathrm{TSZ}}}$, using

the full-sky S10 simulations. We set $\rho^{\kappa \times \mathrm{CIB}}(\ell)$ and $\rho^{\kappa \times \mathrm{TSZ}}(\ell)$ to 0.7 and -0.5 following the discussion above. Finally, we generate a Gaussian random field, M_G, such that $C_\ell^G = \langle M_G, M_G \rangle = C_\ell^{\kappa\kappa} - C_\ell^{\kappa \times FG}$, where we define $C_\ell^{\kappa \times FG}$ as the cross spectra between κ and the combined TSZ plus CIB foregrounds; M_G serves as the part of the κ map uncorrelated with the foregrounds. In the limiting case where we have only one correlated foreground (e.g., let $M_{\mathrm{TSZ}}^{S10} = 0$), Eq. (1) simplifies to Eq. (3.2) from [70]. In this simplified case, it is apparent that Eq. (1) is constructed to

give the correct cross power spectrum between κ and the correlated foreground.

Since we are constructing κ as a scaled combination of M_{CIB}^{S10} and M_{tsz}^{S10}, we include the maximum amount of non-Gaussian features from these foregrounds in our κ maps. Therefore, higher-order correlations between these new κ maps and the foreground components are expected to be larger than those found in the original S10κ simulations. In addition, a higher correlation between κ and foregrounds typically leads to a more significant foreground bias. In these respects, we expect the bias we find with these new κ maps to be conservative.

Since there is a small difference between the power spectra of the CIB and TSZ patch maps (M_{CIB}^{S10} and M_{TSZ}^{S10}) and the full-sky average spectra used in Eq. (1), we find that the new κ maps have a slight excess of power compared to the expected theory power spectrum, $C_\ell^{\kappa\kappa}$. Thus, for each of the four 400 square degree patches, we compute a transfer function by taking the ratio of the average power spectrum of 30 independent κ simulations and the CAMB theory $C_\ell^{\kappa\kappa}$. We find that the transfer function is less than 1 percent of the theory power between $\ell = 2000$ and $\ell = 28000$. We apply this transfer function to the final version of κ maps used in this analysis. In the Appendix A, we show an image of one of our non-Gaussian κ maps, and compare it to a Gaussian κ map generated with the same $C_\ell^{\kappa\kappa}$. We generate 25 realizations of these κ maps for each of the 400 square degree patches. With these new κ maps in hand, we generate lensed CMB maps by first making 100 Gaussian unlensed CMB realizations of the same footprint as the foregrounds, using a theory spectra generated with *Planck2018* cosmological parameters [71]. Then, we lens the unlensed CMB maps with the new κ maps using the publicly available *pixell* library.[1]

B. Creation of foreground-reduced maps

We also make foreground-reduced versions of the CIB, TSZ, and radio patch maps, with flux and cluster cuts matching what we expect to be achieved by CMB-HD. CMB-HD is expected to have a resolution of 0.25 and 0.42 arcminutes, and a sensitivity of 0.8 and 0.7 μK-arcmin white noise, for the 150 and 90 GHz frequency channels, respectively [49,50]. Given these noise levels and resolution, we expect to detect point sources (i.e., galaxies) above 0.04 mJy at 90 GHz with more than 5σ significance [72]. This detection threshold assumes the absence of source confusion from blended sources, which is a reasonable assumption for radio sources since they are not so densely clustered; for reference, we expect the number density of radio sources to be less than 0.07 radio source per 0.42 arcminute beam and 0.03 radio source per 0.25 arcmin beam [59]. The 5σ detection thresholds given in [72] are

based on an analytic calculation assuming only white noise, which is the dominant source of noise at the small scales considered here. Thus we assume that radio sources can be detected and removed that are above 0.04 mJy at 90 GHz, and correspondingly 0.03 mJy at 150 GHz, assuming a spectral index of -0.8. (We note that there will be some scatter in the extrapolation of radio fluxes from 90 to 150 GHz for a source too faint to be detected in the 150 GHz maps directly, and we ignore this scatter in this analysis.) We assume that the point source removal is done via template subtraction, since the point sources have a known shape corresponding to the instrument beam and their fluxes can be measured; the advantage of template subtraction is that we do not cut holes in the map or disturb the coincident lensing signal.

For CIB sources, given CMB-HD noise levels and resolution, as well as the confusion of blended CIB sources (which are more numerous and clustered than radio sources), we find that sources above 0.2 mJy can be detected at the 5σ level in 150 GHz maps directly [73]. To get to a lower flux threshold in 90 and 150 GHz maps, we exploit the fact that CMB-HD will have a 280 GHz channel with 2.7 μK-arcmin white noise. We find that CIB sources above 0.15 mJy at 280 GHz can be detected at the 5σ level in the 280 GHz channel; this detection level was determined by applying a matched filter on maps that included confusion from other CIB sources, as well as the KSZ, TSZ, and CMB [73]. Assuming a spectral index of 2.6 for CIB sources [59], this results in the identification of sources above 0.03 mJy at 150 GHz and 0.008 mJy at 90 GHz. Once detected, we assume these sources can be removed using template subtraction, using a fluxes extrapolated from the 280 GHz map. Since there can be uncertainty on the CIB spectral index for a given source, this can lead to over- or undersubtraction of CIB sources at 90 and 150 GHz. Thus in our analysis, we incorporate some amount of mis-subtraction (see below); we find that this mis-subtraction of CIB sources has minimal impact on the lensing bias (see Sec. IV for more details).

To create the foreground-reduced point source maps, with the flux cuts discussed above, we remove sources in the CIB and radio maps prior to adding them to the lensed CMB maps. To do this, we find all the pixel locations of CIB sources with fluxes above 0.03 mJy at 150 GHz, and set them to the mean flux of the remaining CIB map (which includes only sources below 0.03 mJy); this mimics the diffuse CIB that will still reside at the template-subtracted locations. We follow a similar procedure for the 90 GHz map. To account for some uncertainty in the CIB spectral index for sources whose flux is extrapolated from 280 GHz, we "mis-subtract" sources with fluxes between 0.03 and 0.2 mJy at 150 GHz by up to 5% of their original intensity in both the 150 and 90 GHz maps. For the sources in this flux range, we add back to the appropriate locations their original pixel values multiplied by random numbers drawn

---

[1]https://github.com/simonsobs/pixell.

from $\mathcal{U}[-0.05, 0.05]$. To remove the radio sources, we zero out all radio sources with fluxes above 0.04 mJy at 90 GHz from both the 90 and 150 GHz maps. We perform these cuts in the S10 patches at 0.5 arcminute resolution, before projecting these maps to 0.25 arcminute resolution; we do this because the S10 model was only matched to observations in its native resolution of 0.5 arcminute, and may not necessarily match observations if we extrapolate it to higher resolutions than it was intended to model.[2] Since the radio sources are not correlated with the lensing potential map by construction, after obtaining the power spectra of the residual radio sources at each frequency, we model this component as a Gaussian random field.

We also assume the removal of all 3σ detected TSZ clusters. These clusters will be found based on their integrated cluster Compton-y parameter, Y, which is closely related to cluster mass. We use the cluster mass threshold for 5σ detected clusters found for CMB-HD by [74] and shown in their Fig. 3, and we convert this to a mass threshold for 3σ detected clusters by multiplying their mass threshold by 0.75, which assumes their Y-M scaling relation of $Y \propto M^{1.79}$. We assume clusters above this mass threshold are perfectly subtracted from our maps. To simulate the TSZ cluster subtraction, we find the coordinates of the TSZ clusters with masses above the mass threshold using the S10 cluster catalog. We then mask these TSZ clusters in the TSZ maps by replacing them with 5 and 10 arcminute radius holes for the TSZ clusters with $z > 0.5$ and $z < 0.5$, respectively. Since a few of the closest TSZ clusters, which have irregular shapes and large angular diameters (≈ 1 degree), are not fully covered by the TSZ masks, we visually find these irregular sources in the TSZ maps and set their pixel values to zeros. Note that the masking is done directly in the TSZ maps, simulating a near-perfect subtraction. We discuss mis-subtraction in the context of CIB in Sec. IV and leave an in-depth analysis of cluster mis-subtraction to future work.

The KSZ signal consists of two components, (1) a contribution from the epoch of reionization and (2) a late-time, lower-redshift, component (hereafter reionization KSZ and late-time KSZ, respectively) [75]. We assume that the late-time KSZ can be removed from the large-scale gradient map with an overlapping galaxy survey [75,76], such as from the Rubin Observatory [77], and that it can be removed from the small-scale map with machine learning techniques (e.g., [66,78–83]) (see Sec. V). The reionization KSZ originates at early times ($z \approx 10$) and is more Gaussian than

the lower redshift foregrounds. In addition, the CMB lensing kernel peaks at $z \approx 2$ and thus has limited overlap with the reionization KSZ. Previous work focusing on large scales has found the lensing bias from the reionization KSZ to be negligible for a CMB-S4-type experiment [84]; we leave exploration of the lensing bias from the reionization KSZ for a CMB-HD-type experiment to future work. Here, we approximate the reionization KSZ as a Gaussian random field, which adds noise to the lensing reconstruction but not bias.

In Fig. 1, we show the power spectra of our simulations at 150 (left) and 90 GHz (right). The theoretical CMB power spectra are calculated using CAMB [69] with *Planck2018* cosmological parameters [71]. The instrumental noise spectra are computed following the CMB-HD technical requirements [50]. The power spectra of the foregrounds reflect the reduced-foreground model described above.

III. METHOD

A standard way to measure a lensing signal from CMB maps is to reconstruct a lensing convergence map, κ, using a quadratic estimator ($\mathcal{Q}$) [85–89]. In the presence of non-Gaussian foregrounds correlated with a κ map, the reconstructed lensing auto spectrum can be decomposed into four terms: (1) the lensing signal, (2) the primary bispectrum bias, (2) the secondary bispectrum bias, and (4) the trispectrum bias [51–56]. Modeling the signal (S) as the sum of the lensed CMB temperature map (T) and the extragalactic foregrounds (FG) (i.e., $S = T + FG$), we have

$$\underbrace{\langle \mathcal{Q}[S,S], \mathcal{Q}[S,S] \rangle}_{\text{lensing auto spectrum}} = \underbrace{\langle \mathcal{Q}[T,T], \mathcal{Q}[T,T] \rangle}_{\text{lensing signal}}$$

$$+ \underbrace{2\langle \mathcal{Q}[T,T], \mathcal{Q}[FG,FG] \rangle}_{\text{primary bispectrum bias}}$$

$$+ \underbrace{4\langle \mathcal{Q}[T,FG], \mathcal{Q}[T,FG] \rangle}_{\text{secondary bispectrum bias}}$$

$$+ \underbrace{\langle \mathcal{Q}[FG,FG], \mathcal{Q}[FG,FG] \rangle}_{\text{trispectrum bias}}. \quad (2)$$

In this work, we use the quadratic estimator described in [90] (hereafter HDV estimator), which is tailored to reconstruct κ maps on small scales. On small scales, lensing induces small dipole perturbations that are aligned with the smoothly varying background primordial CMB gradient [91–94]. The HDV estimator looks for these perturbations aligned with the background gradient, by taking as input two versions of filtered CMB maps. One map is filtered to keep only large scales, isolating the smoothly varying background CMB gradient, and we refer to this as the "large-scale gradient map." The other map is filtered to keep only small scales, isolating the perturbations, and we call this the "small-scale map."

---

[2]The S10 simulations include a catalog of CIB sources in which many low flux sources add together to a larger flux source when placed in a 0.5 square arcminute pixel. If we had instead populated 0.25 arcminute maps with the CIB sources from the catalog, then we would have obtained a higher CIB power spectrum than observations allow. Thus we make the CIB flux cut in maps at the native S10 resolution of 0.5 arcminute, and then project the residual sources to higher resolution by upsampling the maps.
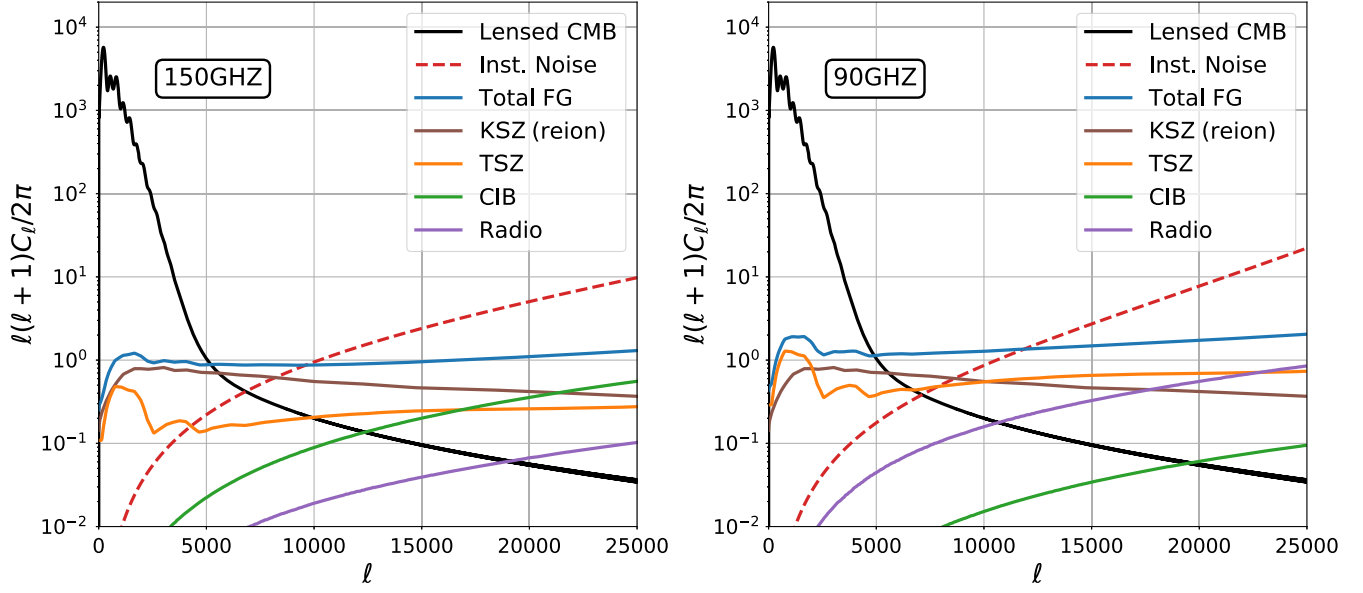
FIG. 1. Shown are the expected power spectra at 150 GHz (left) and 90 GHz (right) for a CMB-HD survey. The theoretical CMB power spectra are calculated using CAMB [69] with *Planck2018* cosmological parameters [71]. Instrumental noise curves are shown as red dashed curves. CMB-HD is expected to have a resolution of 0.25 and 0.42 arcminutes, and a sensitivity of 0.8 and 0.7 μK-arcmin noise, in the 150 and 90 GHz channels, respectively [50]. The reionization KSZ spectrum is shown as the brown curve [75]. The TSZ, CIB, and radio power spectra are computed using the S10 simulations, with the scalings and source cuts as described in Sec. II. In particular, we assume CIB sources above 0.03 and 0.008 mJy at 150 and 90 GHz, respectively, are removed, as well as radio sources above 0.03 and 0.04 mJy at 150 and 90 GHz, respectively. Note that we also account for some amount of CIB mis-subtraction due to uncertainty in the CIB spectral index for source fluxes extrapolated from higher frequencies, as discussed in Sec. II. We also assume that TSZ clusters detected at the 3σ level are removed from the maps. The total foreground spectra (blue curves) are the sum of the reionization KSZ, TSZ, CIB, and radio spectra.

To minimize foreground bias, we only include CMB multipoles in the range of $\ell \in [5000, 30000]$ in the small-scale map. Filtering out modes with $\ell < 5000$ in the "small-scale leg" reduces contributions to the bias from the TSZ signal, which peaks at $\ell \approx 3000$, and the clustering of galaxies, with negligible loss in SNR [47]. For the large-scale gradient map, we only include CMB multipoles in the range of $\ell \in [200, 2000]$. In addition, foreground components with known frequency dependence, such as the TSZ, CIB, and radio, can be removed by leveraging multiple frequency channels; the constrained internal linear combination technique can deproject these foregrounds from a CMB map with some increase in noise [56,95]. However, the deprojection noise penalty is expected to be negligible for the signal dominated large-scale gradient map [57]. Thus in this work we assume that we can completely remove the TSZ and CIB from the gradient map with negligible increase in noise.[3]

Since we assume that we can remove non-Gaussian foregrounds in the large-scale gradient map perfectly, that eliminates the primary bispectrum bias and the trispectrum bias terms in Eq. (2). This simplifies Eq. (2) as follows[4]:

$$\underbrace{\langle \mathcal{Q}[S_{noFG}^{\text{large}}, S^{\text{small}}], \mathcal{Q}[S_{noFG}^{\text{large}}, S^{\text{small}}]\rangle}_{\text{lensing auto spectrum}}$$

$$= \underbrace{\langle \mathcal{Q}[T^{\text{large}}, T^{\text{small}}], \mathcal{Q}[T^{\text{large}}, T^{\text{small}}]\rangle}_{\text{lensing signal}}$$

$$+ \underbrace{\langle \mathcal{Q}[T^{\text{large}}, FG^{\text{small}}], \mathcal{Q}[T^{\text{large}}, FG^{\text{small}}]\rangle}_{\text{secondary bispectrum bias}}, \qquad (3)$$

where the only bias term left is the secondary bispectrum bias.

We note that a "gradient-inversion" (GI) estimator can also reconstruct the lensing potential on small scales [96,97]. In the small-scale, low-noise limit, the GI method should yield even lower lensing reconstruction noise than a QE estimator. We also expect a similar set of foreground mitigation techniques to apply to the GI method; however,

---

[3]This deprojection of the large-scale map can be achieved with upcoming Simons Observatory (SO) data [60] with an increase in noise that is below the sample variance uncertainty, and thus negligible. The removal of the TSZ will be nearly perfect given the known frequency dependence, however, there may be some residual CIB and radio signal due to uncertainty in the spectral indices. We discuss this further in Sec. V.

[4]The factor of 4 in front of the secondary bispectrum bias in Eq. (2) disappears in Eq. (3) because the HDV estimator is not symmetric in the two legs.

since these techniques have been less explored in the context of the GI method, we use the HDV estimator throughout this work.

A. Lensing reconstruction

We perform lensing reconstructions on maps that include the lensed CMB temperature and the foreground-reduced maps discussed in Sec. II. The foreground-reduced maps include the non-Gaussian CIB and TSZ components, correlated with the κ map and each other, as well as the radio and reionization KSZ components modeled as Gaussian random fields. We also add the instrumental white noise levels for CMB-HD, and show the power spectra of these components for 90 and 150 GHz in Fig. 1 as discussed above.

We coadd the 90 (M_{90}) and 150 GHz (M_{150}) maps by the inverse of their noise spectra, $N(\ell)_i$ for $i \in [90 \text{ GHz}, 150 \text{ GHz}]$, such that

$$M_{\text{coadd}} = \frac{W_{90}}{W_{\text{tot}}} M_{90} + \frac{W_{150}}{W_{\text{tot}}} M_{150}, \quad (4)$$

where $W_i = 1/N(\ell)_i$ and $W_{\text{tot}} = \Sigma_i W_i$. We use the publicly available *symlens* library[5] to do the lensing reconstructions with and without the non-Gaussian foregrounds added to the maps. We include the Gaussian radio signal and the reionization KSZ to both the "with" and "without" non-Gaussian foreground cases; we do not expect either of these to introduce a bias, but we do this to keep consistency between the two cases. We also include the radio and reionization KSZ signals in both the large-scale and small-scale maps, anticipating that removing the reionization KSZ will be challenging in general at all scales, and allowing for the increased noise from unresolved radio sources.

We perform lensing reconstructions in the range of $L \in [5000, 30000]$ using only the lensed CMB temperature maps; at these small scales lensing reconstructions from temperature maps, as opposed to polarization maps, dominate the SNR [47]. As mentioned above, for the large-scale gradient leg in the HDV estimator, we include CMB multipoles in the range of $\ell \in [200, 2000]$. For the small-scale leg, we use CMB multipoles in the range of $\ell \in [5000, 30000]$. We process the simulations described in Sec. II to obtain 100 simulated lensing reconstructions.

B. Lensing auto spectra

After generating the lensing reconstructions, we take their power spectra. Since we mask the simulations by a cosine window function (M^{win}) before making the reconstructions, we divide the resulting lensing auto spectra by $w_4 = \sum_j (M_j^{\text{win}})^4 / \sum_j 1$, where j sums over all the pixels

[5]https://github.com/simonsobs/symlens.

in M^{win}; this accounts for the misnormalization of the auto spectra due to the window function. We then subtract the realization dependent lensing bias (RDN0), which arises from the Gaussian component of the auto spectrum [98], and the N1 bias, which arises from higher-order corrections to the auto spectrum [99], from each of the 100 reconstructed auto spectra following [87,100]. We discuss in detail our RDN0 and N1 computations in Appendix B. CMB lensing analyses often further subtract a mean-field κ map, which is a transfer function that mainly affects large scales ($L < 100$), from reconstructed κ maps; however, since the mean field contribution is negligible compared to the reconstruction noise in the range of $L \in [5000, 30000]$, it can be safely ignored.

While the κ realizations (and hence also their power spectra) are not fully independent from each other due to the shared foreground components (M_{CIB}^{S10} and M_{TSZ}^{S10}), these simulated spectra are sufficient to compute a covariance matrix. On the scales of $L \in [5000, 30000]$, the lensing reconstruction is noise dominated. Therefore, the sample variance contribution to the total covariance matrix is subdominant to that from the reconstruction noise, N_L. Thus, we compute a covariance matrix using the 100 simulated lensing power spectra after subtracting the RDN0 and the N1 biases from each simulation.

Since the simulations have a sky area of 400 square degrees, we rescale the simulated covariance matrix to match the expected covariance for an observation over half the sky, $f_{\text{sky}} = 0.5$. While the RDN0 subtraction is important in making the covariance matrix roughly diagonal [47], the off-diagonal components are not negligible. We find that the off-diagonal components of our simulation-based covariance matrix are not fully converged after processing 100 simulations. Thus, we replace the off-diagonal elements of our covariance matrix with those computed by [47], which work used 1000 simulations and a lensing estimator and experimental configuration (0.25 arcmin resolution and 0.5 μK-arcmin white noise) nearly identical to that used in our analysis. In addition, that work included the reionization KSZ in the covariance matrix calculation, however, no additional foregrounds. Given that the TSZ and CIB foregrounds contribute only 20% to the diagonals of our covariance matrix, we do not expect them to appreciably contribute to the off-diagonal elements. In Appendix C, we compare a covariance matrix obtained from an analytic calculation with our simulation-based covariance matrix and discuss the impact of the off-diagonal elements of the covariance matrix on the SNR.

To calculate the bias to the lensing signal from the non-Gaussian correlated TSZ and CIB foregrounds, we compute the lensing auto spectra with and without them, which we label $C_L^{\kappa\kappa,\text{NGFG}}$ and $C_L^{\kappa\kappa,\text{GFG}}$, respectively. Then, we compute the foreground bias, $C_L^{FG\,\text{bias}}$, as the difference between the two lensing auto spectra. For the case "without" these correlated non-Gaussian foregrounds (i.e.,

$C_L^{\kappa\kappa,\text{GFG}}$), we match the 2D power spectrum of the non-Gaussian foregrounds patch by patch. To do this, we generate Gaussian TSZ and CIB maps by randomizing the phases of the non-Gaussian TSZ and CIB maps; this phase randomization not only makes the resulting TSZ and CIB maps Gaussian but also breaks the correlations between TSZ, CIB, and κ. We further improve the match between these two cases by using the same realizations of CMB, κ, and instrumental noise. These shared realizations, plus the phase-randomized foreground maps, make the RDN0 and N1 biases closely match for the two cases, minimizing any apparent lensing bias due to slight mis-estimations of RDN0 and N1. Explicitly the lensing auto spectrum bias is given by

$$
\begin{aligned}
C_L^{FG\,\text{bias}} &= \frac{1}{100}\sum_{m=1}^{100}(C_{L,m}^{\kappa\kappa,\text{NGFG}} - C_{L,m}^{\kappa\kappa,\text{GFG}}), \\
&= \frac{1}{100}\sum_{m=1}^{100}[(C_{L,\text{raw},m}^{\kappa\kappa,\text{NGFG}} - RDN_{0,m}^{\text{NGFG}} - N_1^{\text{NGFG}}) \\
&\quad - (C_{L,\text{raw},m}^{\kappa\kappa,\text{GFG}} - RDN_{0,m}^{\text{GFG}} - N_1^{\text{GFG}})],
\end{aligned}
\tag{5}
$$

where m sums over the 100 reconstructed auto spectra. Note that the N1 bias term does not have the m subscript because the same N1 bias is subtracted from each simulation. Here, any additive bias not induced by the non-Gaussian correlated TSZ and CIB maps will cancel out in the difference. Similarly, any multiplicative bias will cancel out in the ratios used to calculate the fractional bias or SNR, since the denominators are also a function of the simulated auto spectra.

IV. RESULTS

Figure 2 shows the effect of the extragalactic foregrounds on the lensing reconstruction in the range of $L \in [5000, 20000]$; this L range is chosen because it contributes most of the SNR in determining if there is a problem with CDM on small scales [47]. Following [47], we use a uniform binning of $\Delta\ell = 300$. The left panel shows the fractional lensing bias for various foreground scenarios (CIB alone, TSZ alone, and CIB plus TSZ). Here, the foreground bias is computed as the difference between the reconstructed lensing auto spectrum with and without the non-Gaussian TSZ and/or CIB foregrounds. All cases have the same realizations of CMB, κ, instrumental noise, and Gaussian radio and reionization KSZ foregrounds, such that the difference in the lensing spectra is solely due to the TSZ/CIB foregrounds. The total foreground bias (blue curve) includes both CIB and TSZ, and is within a few percent of the unbiased spectra over the range of interest.
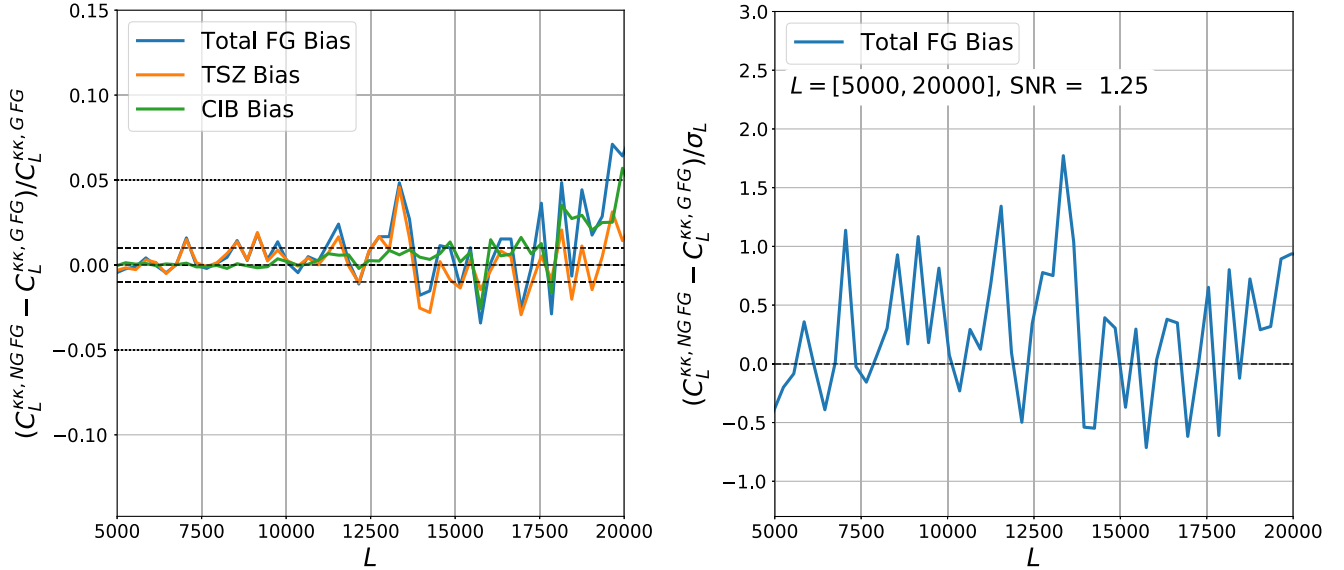

FIG. 2. Left: shown is the fractional lensing bias for various foreground scenarios. The total foreground case (blue curve) includes both TSZ and CIB foregrounds, while the other two cases include either TSZ alone or CIB alone. Here, TSZ and CIB foregrounds are either modeled as uncorrelated Gaussian random fields (G) or prepared from the non-Gaussian (NG) simulations discussed in Sec. II. We include the reionization KSZ and radio sources, modeled as Gaussian random fields, to all scenarios. The black dashed lines indicate 1% bias, while the black dotted lines indicate 5% bias. We find that the foreground bias is below a few percent for most of the range $L \in [5000, 20000]$. Right: shown is the SNR of the total foreground bias (both TSZ and CIB) over half the sky, $f_{\text{sky}} = 0.5$. The error bars are the square root of the diagonal elements of the covariance matrix described in Sec. IV. The simulation-based covariance matrix is computed by processing 100 simulations described in Sec. II through the lensing reconstruction and auto spectrum pipeline described in Sec. III. Using the full covariance matrix, we find that the total foreground bias to the lensing auto spectrum deviates from zero by only 1.25σ in the range of $L \in [4800, 20100]$.

The right panel of Fig. 2 shows the total foreground bias divided by the diagonal error bars of the simulation-based $C_L^{\kappa\kappa}$ covariance matrix discussed in Sec. III. We find that the total foreground bias is below 0.5σ for many spectral bins in the range shown. In order to determine whether this foreground bias is detectable, we use the full covariance matrix to calculate the SNR as follows. The χ^2 of the total foreground bias with respective to a null signal (i.e., $C_L^{\rm bias} = 0$) is computed as

$$\chi_{\rm null}^2 = \sqrt{\sum_{L,L'}(X_L - Y_L)C_{L,L'}^{-1}(X_{L'} - Y_{L'})}, \qquad (6)$$

where $X_L = C_L^{\kappa\kappa,\rm NGFG} - C_L^{\kappa\kappa,\rm GFG}$, $Y_L = 0$, and $C_{L,L'}^{-1}$ are the elements of the inverted covariance matrix. We calculate this over the range of $L \in [4800, 20100]$. We compute $\chi_{\rm best-fit}^2$ by replacing $Y_L = 0$ with the best-fit $C_L^{\rm bias}$ line obtained from fitting a straight line in the range of $L \in [4800, 20100]$; for the total FG case we find a slope of 7.91×10^{-17} and an intercept of 2.74×10^{-14} for this best-fit. We also obtain best-fit lines for the TSZ-only and CIB-only cases. We find that the lensing bias in the total FG case favors the best-fit line over the null at only 1.25σ, where we compute this significance as $\sqrt{\chi_{\rm null}^2 - \chi_{\rm best-fit}^2}$. We find that the significance of detecting the foreground bias for the TSZ-only and CIB-only cases is 0.59σ and 1.11σ, respectively. As discussed in more detail in Appendix C, the inclusion of the off-diagonal elements of the covariance matrix is important for the total SNR calculation. We find that the total SNR would increase by a factor of 2 if we used only the diagonal elements to compute the SNR.

Figure 3 shows the SNR for distinguishing between an $m = 10^{-22}$ eV FDM model [25,30] and a CDM model, in the presence of extragalactic foregrounds. Here, the signal is the difference between a theoretical $m = 10^{-22}$ eV FDM $C_\ell^{\kappa\kappa}$ and a theoretical CDM $C_\ell^{\kappa\kappa}$ obtained from the WarmAndFuzzy code [101] and used in [47]. Here we use the same error bars and covariance matrix used in the right panel of Fig. 2. We compute the total SNR in the range of $L \in [4800, 20100]$ following

$$\rm SNR = \sqrt{\sum_{L,L'}(X_L - Y_L)C_{L,L'}^{-1}(X_{L'} - Y_{L'})}, \qquad (7)$$

where $X_L = C_L^{\kappa\kappa,\rm FDM}$, $Y_L = C_L^{\kappa\kappa,\rm CDM}$ (the FDM and CDM $C_L^{\kappa\kappa}$, respectively), and $C_{L,L'}^{-1}$ are the elements of the inverted covariance matrix, including off-diagonal elements. Calculating this, we obtain a SNR of 5.06σ. As shown in the right panel of Fig. 2, the foreground bias in the same L range is not detectable.
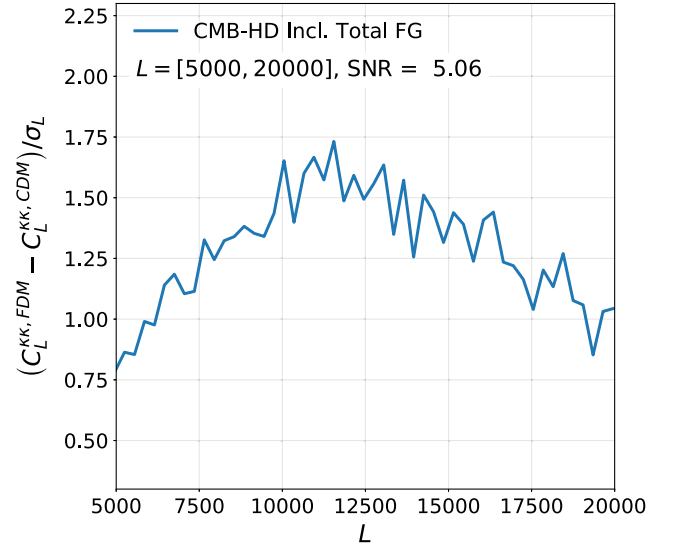

FIG. 3. Shown is the SNR for distinguishing an $m = 10^{-22}$ eV FDM model from a CDM model when including the noise from extragalactic foregrounds in the covariance matrix. The foregrounds include the reionization KSZ, TSZ, CIB, and radio sources, with the source cuts discussed in Sec. II, resulting in the power spectra shown in Fig. 1. Here, the signal is the difference between a theoretical 10^{-22} eV FDM $C_\ell^{\kappa\kappa}$ and a theoretical CDM $C_\ell^{\kappa\kappa}$ obtained from the WarmAndFuzzy code [101]. We do not include bias from the foregrounds in either the CDM or FDM models here. The error bars are the square root of the diagonals of the covariance matrix described in Sec. III. Using the full covariance matrix, we find a SNR of 5.06σ in the range of $L \in [4800, 20100]$. As shown in the right panel of Fig. 2, the foreground bias is not detectable in this same L range.

A. Uncertainty in the CIB spectral index

Even though CMB-HD will detect CIB sources with an unprecedented sensitivity [48–50], there will likely be some uncertainty in the CIB spectral indices for the fainter sources. Since we template-subtract some dimmer CIB sources by extrapolating their flux measurements at the 280 GHz channel down to 150 and 90 GHz, this uncertainty can lead to a mis-subtraction of the CIB sources. In our baseline analysis, we include this mis-substraction effect (see Sec. II for details). Here, we quantify its effect on the foreground bias and the simulation-based covariance matrix.

In our baseline analysis, we "mis-subtract" CIB sources in the flux range of 0.03 to 0.2 mJy at 150 GHz, in both 90 and 150 GHz maps by up to 5%. These sources will be detected in 280 GHz maps with a significance between 5 and about 30σ, and thus will have some uncertainty from the flux measurement itself at 280 GHz as well as some uncertainty in the CIB spectral index when extrapolating to 150 GHz. We pick a mis-subtraction level up to 5% for all the sources, noting that for the bright sources, 5% mis-subtraction may be conservative since their fluxes are measured with better precision than that, and they will

also be detected at 220 GHz, thereby lowering the uncertainty in their spectral index; for fainter sources, a larger mis-subtraction amount may be warranted, however, since these sources are faint, that should minimally increase the noise, and, as we show below, may serve to lower the lensing auto spectrum bias.

We also run a case where these CIB sources are subtracted perfectly and their pixel values replaced with the mean flux of the residual CIB sources. Figure 4 shows the effect of this CIB mis-subtraction compared to the perfect subtraction case. Overall, the lensing bias from the CIB component stays roughly the same in the range of $L \in [4800, 20100]$. When including the effect of mis-subtracting the CIB sources, the SNR of the total FG bias goes from 1.73σ to 1.25σ. The reduction in the SNR is likely due to both the increase in the lensing reconstruction noise and the decrease in the CIB foreground bias. For the latter, the random "mis-subtraction" may break some correlation between the κ and CIB maps resulting in the lower foreground bias. The effect of the mis-subtraction on distinguishing a 10^{-22} eV FDM model from a CDM model (see Fig. 3) is also negligible. In this case, the total SNR
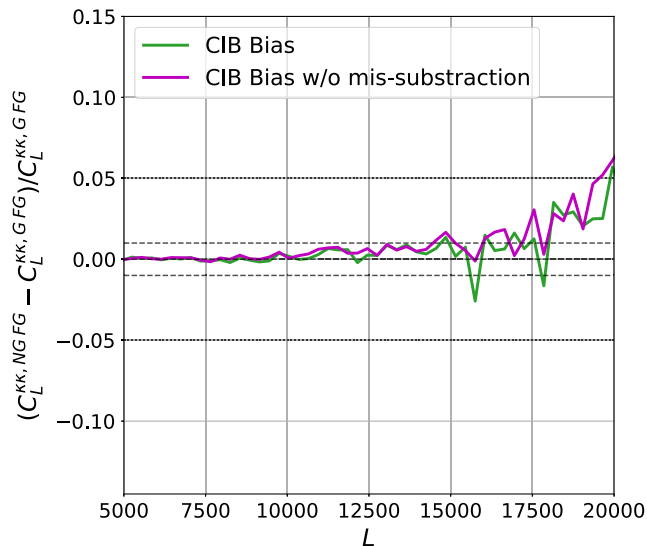
goes from 5.11σ to 5.06σ when we include CIB mis-subtraction in our simulations.

V. DISCUSSION

The results above suggest that the two most dominant foregrounds, the TSZ and the CIB, can be reduced to a level that will not significantly bias the lensing auto spectrum in the range $L \in [4800, 20100]$. In addition, the noise penalty in using the HDV estimator and including the residual foregrounds still allows for the ability to determine if CDM is a good match to the small-scale matter power spectrum, or if structure is suppressed at the level, for example, that a 10^{-22} eV FDM model would predict.

Below we outline some additional foreground contributions that were not included in this analysis, and we outline how we foresee them being mitigated. In future work, these additional contributions will be explored in more detail.

A. Foregrounds in the large-scale gradient map

In the current analysis, we assumed that there were no TSZ and CIB foregrounds in the large-scale gradient map. For this map, we only included CMB multipoles in the range of $\ell \in [200, 2000]$. We anticipate that these large-scale gradient maps will be provided by Simons Observatory data, which extends down to CMB multipoles of $\ell = 30$ [60], as opposed to $\ell_{\min} = 1000$ for CMB-HD [48]. Since the frequency dependence of the TSZ signal is exactly known, we can deproject it using the constrained internal linear combination technique [56,95]. While there is some noise penalty in using this deprojection technique, given the SO noise levels, the temperature maps will still be sample variance limited over half the sky for $\ell \in (200, 2000)$; the increase in noise from the deprojection will be negligible compared to the uncertainty from the sample variance [57].

CIB and radio sources can be similarly deprojected from the large-scale gradient map. However, uncertainty in the spectral indecies will leave some residual foregrounds. A path forward would be to template-subtract the same sources detected in CMB-HD maps from the SO maps, and then deproject the SO maps using their multi-frequency data. In addition, CCAT-p data, which is designed to map the CIB on large scales, can be used to further clean the CIB from the gradient leg [102]. Given that our current analysis shows that some amount of source mis-subtraction actually lowers the bias to the lensing auto spectrum; we do not expect residual CIB and radio sources in the gradient map to add significantly to the lensing bias. In addition, we can apply point-source hardening techniques to our lensing reconstruction procedure; point-source hardening is shown to reduce the foreground bias from point sources with only small increases in the reconstruction noise [55,56]. We leave a detailed analysis of this to future work.

FIG. 4. Shown is the fractional lensing bias due to the CIB foreground with and without accounting for some uncertainty in the CIB spectral indecies. The black dashed lines indicate 1% bias, while the black dotted lines show 5% bias. The magenta curve represents the case where we can perfectly remove all the CIB sources in the range of 0.03 to 0.2 mJy at 150 GHz. The green curve shows the case where some sources are either over or under subtracted by up to 5% of their original fluxes in both 90 and 150 GHz maps. After including the effect of the mis-subtraction, the SNR of the total foreground bias (including both CIB and TSZ) in the range of $L \in [4800, 20100]$ gets smaller, changing from 1.73σ to 1.25σ (see Sec. IV for details). The total SNR in distinguishing a 10^{-22} eV FDM model from a CDM model changes only slightly from 5.11σ to 5.06σ.

B. Bias from radio galaxies

We did not explicitly include correlated radio sources in this analysis since we focused on the bias from the TSZ and CIB foregrounds, anticipating these to be the dominant sources of bias as is the case for the lensing auto spectrum on large scales (see e.g., [66]). As mentioned in Sec. II, we expect to be able to template-subtract radio sources above 0.04 mJy in 90 GHz maps at the 5σ level. [Radio sources are subdominant in 150 GHz maps compared to the CIB (see Fig. 1).] Given this low flux cut, and the fact that radio sources are fewer in number and less correlated with the κ map, we do not anticipate significant bias to the lensing auto spectrum from radio sources; however, we leave a detailed investigation to future work.

C. Bias from the late-time KSZ effect

The KSZ effect arises when CMB photons inverse-Compton scatter off ionized gas that has a bulk motion. As discussed in Sec. II, the KSZ effect has two components, (1) one from the epoch of reionization and (2) one from the lower redshift Universe (e.g., [75]). We assume in this work that the reionization KSZ, which is from $z \approx 10$, is Gaussian and uncorrelated with the κ map, whose kernel peaks at much lower redshift ($z \approx 2$). We include the reionization KSZ in both the large-scale gradient map and the small-scale map, to add it as an extra noise source; from Fig. 1, we see that this is a significant source of noise below $\ell = 10000$. In terms of bias to the lensing auto spectrum, the reionization KSZ has been shown to lead to negligible bias on larger scales than studied here [84], and future investigation is warranted to explore its effect on smaller scales.

The late-time KSZ may lead to considerable bias to the lensing auto spectrum. However, there are a number of promising avenues for removing it. For the large-scale gradient map, removing the late-time KSZ with an overlapping galaxy survey (e.g., the Vera C. Rubin Observatory [77]) using variations of the techniques discussed in [75,76,103,104] seems promising. For the small-scale map, recent advances in Deep Learning pattern recognition is a promising avenue to remove the late-time KSZ. While the KSZ effect has no frequency dependence, making it hard to remove with multifrequency techniques [53], the KSZ signal has a characteristic spatial pattern that can be recognized by a convolution neural network (e.g., [66,78,80–82]). In turn, we can use this recognized spatial pattern to remove the KSZ signal, especially on small scales above $\ell = 5000$ where it is the dominant signal. Examples of deep learning foreground mitigation methods are shown in [79,83].

Since the KSZ signal is proportional to the bulk motion of the electron gas, and because the motion of the gas toward and away from an observer is equally likely, the mean correlation between the KSZ signal and κ is zero. As a result, we expect the secondary bispectrum bias

[see Eq. (2)], which arises due to the correlation between a non-Gaussian foreground and κ, to be small. If we can clean the late-time KSZ from either the large-scale or small-scale maps, then this secondary bispectrum bias is the only bias term that will remain. We leave a thorough examination of the bias from the late-time KSZ to subsequent work.

VI. CONCLUSION

In this work, we provide a proof of concept demonstration that lensing bias from the two most dominant foregrounds, the TSZ and CIB, can be mitigated, and that CMB-HD can measure the CMB lensing auto spectrum on small scales ($L \in [5000, 20000]$). We find that the combined TSZ and CIB foreground bias, in the coadded CMB-HD 90 and 150 GHz channels, has a SNR of only 1.3σ. We also find that modest mis-subtraction of faint CIB sources due to their uncertain spectral indices has negligible impact. Using an interesting alternative dark matter model as a benchmark, we find that CMB-HD can distinguish a 10^{-22} eV FDM model from a CDM model at the 5σ level. While this analysis provides an initial proof of concept demonstration, a full investigation including all potential sources of bias, and an exploration of the optimal path to mitigate them is warranted. Such measurements of the high-resolution CMB lensing auto spectrum will probe the matter spectrum power spectrum out to small scales ($k \sim 10~h\,\text{Mpc}^{-1}$). In addition to providing a robust measurement of the small-scale matter distribution via gravitational lensing, this will put stringent constraint on alternative dark matter models and baryonic processes, informing both fundamental physics and galaxy evolution.

APPENDIX A: GAUSSIAN AND NON-GAUSSIAN κ MAPS

In Fig. 5, we show a visual comparison between one realization of the non-Gaussian κ map generated via the procedure discussed in Sec. II, and a Gaussian κ map generated with the same $C_\ell^{\kappa\kappa}$. We generate the Gaussian κ map as a Gaussian random field, and use many realizations of them to compute a realization dependent bias (RDN0) and N1 bias (see Sec. III and Appendix B for details). The non-Gaussian κ map is generated using Eq. (1) from Sec. II, and is 70% correlated with the S10 CIB map and 50% correlated with the S10 TSZ map at 150 GHz by
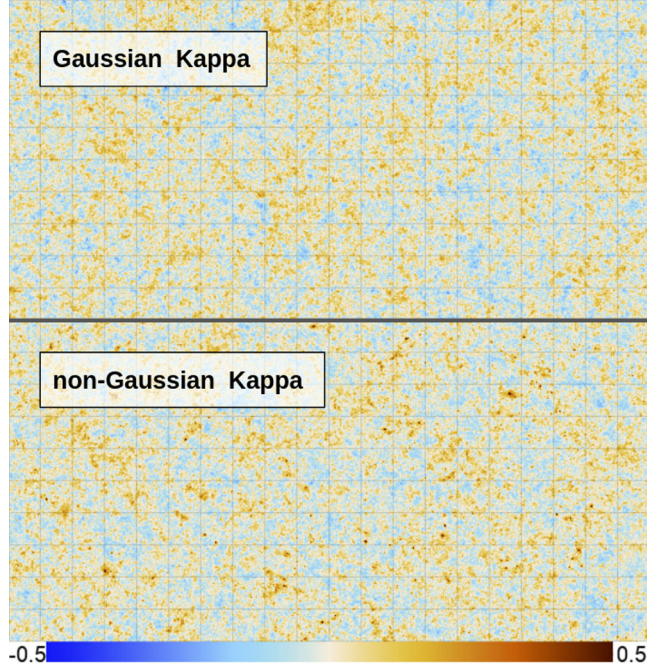
FIG. 5. Shown is a Gaussian (top panel) and a non-Gaussian (bottom panel) κ map. We generate the non-Gaussian κ map using Eq. (1), and the Gaussian κ map as a Gaussian random field with the same $C_\ell^{\kappa\kappa}$. By construction, the non-Gaussian κ map is 70% correlated with the S10 CIB map and 50% correlated with the S10 TSZ map at 150 GHz. Visually, we find that the large-scale fluctuations look similar between the Gaussian and non-Gaussian κ maps, however, there is more small-scale structure in the non-Gaussian κ map that mirrors bright structures in the CIB and TSZ maps. We expect that these κ maps will yield a bias to the lensing auto spectrum that is conservative, as discussed in Sec. II.
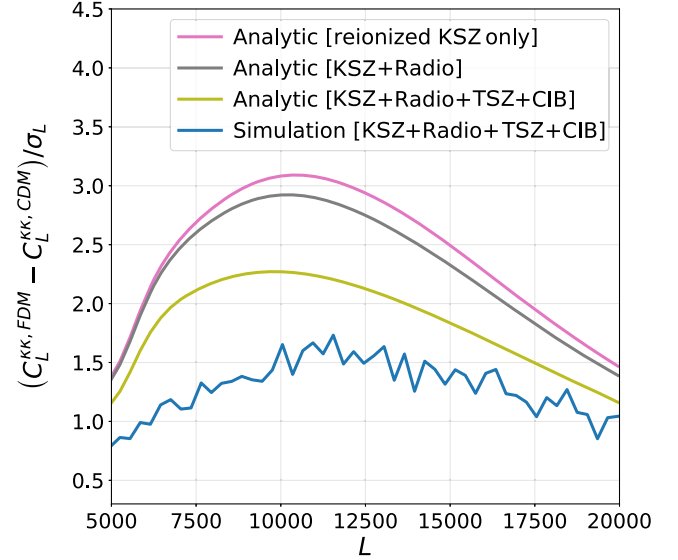

FIG. 6. The same as in Fig. 3, except for three additional curves calculated with analytic error bars. The blue curve is the same as shown in Fig. 3. We compute the analytic error bars for other three curves using Eq. (4) from [105] given the theoretical $C_L^{\kappa\kappa}$ spectrum and analytic N0 lensing noise curves. The same experimental configuration (white noise level, resolution, and sky coverage) is assumed in all four cases. As in [47], we find that the diagonal variance from simulations is larger than that predicted by the analytic calculation. We also find that the off-diagonal variance has a large impact on the total SNR calculation. Using only the diagonal variance, we find that the analytic case (yellow curve) gives a SNR of 13.42σ, while the simulation-based case (blue curve) gives a SNR of 9.37σ. Once we include the off-diagonal variance for the simulation-based case, we find that the SNR reduces by roughly a factor two, going from 9.37σ to 5.06σ.

construction. Note that Eq. (1) leaves in the maximum amount of non-Gaussian structure from the S10 CIB and TSZ maps in the non-Gaussian κ maps. In general, this will lead to larger higher-order correlations between these non-Gaussian foregrounds and the non-Gaussian κ map than expected. Visually, we find that the large-scale fluctuations look similar between the Gaussian and non-Gaussian κ maps. However, we find more small structure in the non-Gaussian κ maps. In particular, the brighter spots in the non-Gaussian κ maps roughly coincide with the location of bright CIB sources.

APPENDIX B: SIMULATION-BASED RDN0 AND N1 COMPUTATIONS

We compute the realization dependent Gaussian lensing bias (RDN0) [98] following Eq. (7) from [87]. Unlike the analytic N0 bias described in [86], the RDN0 bias was designed to account for a slight difference between the power spectra of simulations and data. In this work, our "data" is the realistic non-Gaussian and phase-randomized simulations discussed in Sec. II. Our "simulations" are

Gaussian simulations constructed to match the power spectra of these realistic simulations.

We calculate the RDN0 bias as

$$
\begin{aligned}
C_L^{\mathrm{RDN0}}[XY, AB] = \langle &\mathcal{Q}[XY^S, AB^S] + \mathcal{Q}[XY^S, A^S B] \\
&+ \mathcal{Q}[X^S Y, AB^S] + \mathcal{Q}[X^S Y, A^S B] \\
&- \mathcal{Q}[X^S Y^{S'}, A^S B^{S'}] - \mathcal{Q}[X^S Y^{S'}, A^{S'} B^S] \rangle_{S,S'},
\end{aligned}
$$
(B1)

where $\mathcal{Q}$ stands for a quadratic estimator. Since Eq. (B1) does not assume any symmetry in $\mathcal{Q}$, it works for the HDV estimator without any modification. S and S' superscripts indicate two independent realizations of Gaussian simulations. For each of the 100 "data" simulations, we average over 20 RDN0 computations, where each RDN0 run requires two new Gaussian simulations (for S and S'); thus, in total, we make $2 \times 100 \times 20 = 4000$ Gaussian simulations for the RDN0 calculation. As in [47], we find that subtracting the RDN0 bias removes a large amount

off-diagonal correlation in the simulation-based $C_L^{\kappa\kappa}$ covariance matrix.

Similarly, we compute the N1 bias, which arises from higher-order corrections to the auto spectrum [99], following [87,100]:

$$
\begin{aligned}
C_L^{N1}[XY,AB] = \langle &\mathcal{Q}[X^{S_\phi}Y^{S'_\phi}, A^{S_\phi}B^{S'_\phi}] + \mathcal{Q}[X^{S_\phi}Y^{S'_\phi}, A^{S'_\phi}B^{S_\phi}] \\
&- \mathcal{Q}[X^S Y^{S'}, A^S B^{S'}] - \mathcal{Q}[X^S Y^{S'}, A^{S'}B^S]\rangle_{S,S',S_\phi,S'_\phi}.
\end{aligned}
$$
(B2)

Here, S and S' superscripts again indicate two independent realizations of Gaussian simulations, while S_ϕ and S'_ϕ superscript stand for two independent CMB realizations lensed by the same κ map. Note that the N1 bias is computed using only the Gaussian simulations. We average over 100 N1 computations; using, in total, $4 \times 100 = 400$ Gaussian simulations.

APPENDIX C: ANALYTIC VERSUS SIMULATION-BASED C_L^{KK} VARIANCE

Figure 6 shows the comparison of SNRs for distinguishing a 10^{-22} eV FDM model from a CDM model, computed either using an analytic or simulation-based C_L^{KK} covariance matrix. The error bars for the blue curve are the square root of the diagonal elements of the simulation-based covariance matrix described in Sec. IV. We compute the analytic error bars for the other three curves using Eq. (4) from [105] given the theoretical $C_L^{\kappa\kappa}$ spectrum and analytic N0 lensing noise curves. The analytic N0 lensing noise curves are computed using *symlens* assuming the same experimental configuration (white noise level, resolution, and sky coverage) used in the simulated case. As in [47], we find that the diagonal variance from the simulations is larger than that predicted by the analytic calculation (compare yellow and blue curves). Where the discrepancy between the diagonal variance is largest coincides with where the simulation-based $C_L^{\kappa\kappa}$ band powers are significantly correlated $L \in [5000, 12000]$. In addition, we find that the off-diagonal variance has a large impact on the total SNR. Using only the diagonal of variance, we find that the analytic case (yellow curve) gives a SNR of 13.42σ, whereas the simulation-based case (blue curve) gives a SNR of 9.37σ. However, once we include the off-diagonal variance for the simulation-based case, we find that the SNR reduces by roughly a factor of 2, going from 9.37σ to 5.06σ. Note that the analytic variance calculations do not include any off-diagonal contributions. We also show how the analytic diagonal variance changes with the addition of different foregrounds.

[1] P. J. E. Peebles, Astrophys. J. Lett. **263,** L1 (1982).
[2] G. R. Blumenthal, S. M. Faber, J. R. Primack, and M. J. Rees, Nature (London) **311,** 517 (1984).
[3] M. Davis, G. Efstathiou, C. S. Frenk, and S. D. M. White, Astrophys. J. **292,** 371 (1985).
[4] C. S. Frenk and S. D. M. White, Ann. Phys. (Berlin) **524,** 507 (2012).
[5] J. R. Primack, Ann. Phys. (Berlin) **524,** 535 (2012).
[6] T. M. C. Abbott, M. Aguena, A. Alarcon, S. Allam, O. Alves, A. Amon, F. Andrade-Oliveira, J. Annis, S. Avila et al. (DES Collaboration), Phys. Rev. D **105,** 023520 (2022).
[7] R. A. Flores and J. R. Primack, Astrophys. J. Lett. **427,** L1 (1994).
[8] B. Moore, Nature (London) **370,** 629 (1994).
[9] A. Klypin, A. V. Kravtsov, O. Valenzuela, and F. Prada, Astrophys. J. **522,** 82 (1999).
[10] B. Moore, S. Ghigna, F. Governato, G. Lake, T. Quinn, J. Stadel, and P. Tozzi, Astrophys. J. Lett. **524,** L19 (1999).
[11] J. P. Ostriker and P. Steinhardt, Science **300,** 1909 (2003).
[12] M. Boylan-Kolchin, J. S. Bullock, and M. Kaplinghat, Mon. Not. R. Astron. Soc. **415,** L40 (2011).
[13] A. Brooks, Ann. Phys. (Berlin) **526,** 294 (2014).
[14] D. H. Weinberg, J. S. Bullock, F. Governato, R. Kuzio de Naray, and A. H. G. Peter, Proc. Natl. Acad. Sci. U.S.A. **112,** 12249 (2015).
[15] K. A. Oman, J. F. Navarro, A. Fattahi, C. S. Frenk, T. Sawala, S. D. M. White, R. Bower, R. A. Crain, M. Furlong, M. Schaller et al., Mon. Not. R. Astron. Soc. **452,** 3650 (2015).
[16] A. Del Popolo and M. Le Delliou, Galaxies **5,** 17 (2017).
[17] J. S. Bullock and M. Boylan-Kolchin, Annu. Rev. Astron. Astrophys. **55,** 343 (2017).
[18] P. Colín, V. Avila-Reese, and O. Valenzuela, Astrophys. J. **542,** 622 (2000).
[19] P. Bode, J. P. Ostriker, and N. Turok, Astrophys. J. **556,** 93 (2001).
[20] M. Viel, J. Lesgourgues, M. G. Haehnelt, S. Matarrese, and A. Riotto, Phys. Rev. D **71,** 063534 (2005).
[21] M. S. Turner, Phys. Rev. D **28,** 1243 (1983).
[22] W. H. Press, B. S. Ryden, and D. N. Spergel, Phys. Rev. Lett. **64,** 1084 (1990).
[23] S.-J. Sin, Phys. Rev. D **50,** 3650 (1994).
[24] J. Goodman, New Astron. **5,** 103 (2000).
[25] W. Hu, R. Barkana, and A. Gruzinov, Phys. Rev. Lett. **85,** 1158 (2000).
[26] P. J. E. Peebles, Astrophys. J. Lett. **534,** L127 (2000).

[27] L. Amendola and R. Barbieri, Phys. Lett. B **642**, 192 (2006).

[28] H.-Y. Schive, T. Chiueh, and T. Broadhurst, Nat. Phys. **10**, 496 (2014).

[29] D. J. E. Marsh, Phys. Rep. **643**, 1 (2016).

[30] L. Hui, J. P. Ostriker, S. Tremaine, and E. Witten, Phys. Rev. D **95**, 043541 (2017).

[31] E. D. Carlson, M. E. Machacek, and L. J. Hall, Astrophys. J. **398**, 43 (1992).

[32] D. N. Spergel and P. J. Steinhardt, Phys. Rev. Lett. **84**, 3760 (2000).

[33] M. Vogelsberger, J. Zavala, and A. Loeb, Mon. Not. R. Astron. Soc. **423**, 3740 (2012).

[34] A. B. Fry, F. Governato, A. Pontzen, T. Quinn, M. Tremmel, L. Anderson, H. Menon, A. M. Brooks, and J. Wadsley, Mon. Not. R. Astron. Soc. **452**, 1468 (2015).

[35] O. D. Elbert, J. S. Bullock, S. Garrison-Kimmel, M. Rocha, J. Oñorbe, and A. H. G. Peter, Mon. Not. R. Astron. Soc. **453**, 29 (2015).

[36] M. Kaplinghat, S. Tulin, and H.-B. Yu, Phys. Rev. Lett. **116**, 041302 (2016).

[37] A. Kamada, M. Kaplinghat, A. B. Pace, and H.-B. Yu, Phys. Rev. Lett. **119**, 111102 (2017).

[38] R. Huo, M. Kaplinghat, Z. Pan, and H.-B. Yu, Phys. Lett. B **783**, 76 (2018).

[39] S. Tulin and H.-B. Yu, Phys. Rep. **730**, 1 (2018).

[40] L. Berezhiani and J. Khoury, Phys. Rev. D **92**, 103510 (2015).

[41] J. Khoury, SciPost Phys. Lect. Notes, 42 (2022).

[42] M. P. van Daalen, J. Schaye, C. M. Booth, and C. Dalla Vecchia, Mon. Not. R. Astron. Soc. **415**, 3649 (2011).

[43] A. M. Brooks, M. Kuhlen, A. Zolotov, and D. Hooper, Astrophys. J. **765**, 22 (2013).

[44] A. Natarajan, A. R. Zentner, N. Battaglia, and H. Trac, Phys. Rev. D **90**, 063516 (2014).

[45] A. Schneider, R. Teyssier, J. Stadel, N. E. Chisari, A. M. C. Le Brun, A. Amara, and A. Refregier, J. Cosmol. Astropart. Phys. 03 (2019) 020.

[46] S. K. Giri and A. Schneider, J. Cosmol. Astropart. Phys. 12 (2021) 046.

[47] H. N. Nguyên, N. Sehgal, and M. S. Madhavacheril, Phys. Rev. D **99**, 023502 (2019).

[48] N. Sehgal, S. Aiola, Y. Akrami, K. Basu, M. Boylan-Kolchin, S. Bryan, S. Clesse, F.-Y. Cyr-Racine, L. Di Mascolo, S. Dicker et al., Bull. Am. Astron. Soc. **51**, 6 (2019).

[49] N. Sehgal, H. N. Nguyen, J. Meyers, M. Munchmeyer, T. Mroczkowski, L. Di Mascolo, E. Baxter, F.-Y. Cyr-Racine, M. Madhavacheril, B. Beringue et al., Bull. Am. Astron. Soc. **51**, 43 (2019).

[50] N. Sehgal, S. Aiola, Y. Akrami, K. moni Basu, M. Boylan-Kolchin, S. Bryan, C. M. Casey, S. Clesse, F.-Y. Cyr-Racine, L. Di Mascolo et al., arXiv:2002.12714.

[51] A. van Engelen, S. Bhattacharya, N. Sehgal, G. P. Holder, O. Zahn, and D. Nagai, Astrophys. J. **786**, 13 (2014).

[52] S. J. Osborne, D. Hanson, and O. Doré, J. Cosmol. Astropart. Phys. 03 (2014) 024.

[53] S. Ferraro and J. C. Hill, Phys. Rev. D **97**, 023512 (2018).

[54] E. Schaan and S. Ferraro, Phys. Rev. Lett. **122**, 181301 (2019).

[55] N. Sailer, E. Schaan, and S. Ferraro, Phys. Rev. D **102**, 063517 (2020).

[56] N. Sailer, E. Schaan, S. Ferraro, O. Darwish, and B. Sherwin, Phys. Rev. D **104**, 123514 (2021).

[57] M. S. Madhavacheril and J. C. Hill, Phys. Rev. D **98**, 023534 (2018).

[58] O. Darwish, M. S. Madhavacheril, B. D. Sherwin, S. Aiola, N. Battaglia, J. A. Beall, D. T. Becker, J. R. Bond, E. Calabrese, S. K. Choi et al., Mon. Not. R. Astron. Soc. **500**, 2250 (2020).

[59] N. Sehgal, P. Bode, S. Das, C. Hernandez- Monteagudo, K. Huffenberger, Y.-T. Lin, J. P. Ostriker, and H. Trac, Astrophys. J. **709**, 920 (2010).

[60] P. Ade, J. Aguirre, Z. Ahmed, S. Aiola, A. Ali, D. Alonso, M. A. Alvarez, K. Arnold, P. Ashton, J. Austermann et al., J. Cosmol. Astropart. Phys. 02 (2019) 056.

[61] J. Dunkley, E. Calabrese, J. Sievers, G. E. Addison, N. Battaglia, E. S. Battistelli, J. R. Bond, S. Das, M. J. Devlin, R. Dünner et al., J. Cosmol. Astropart. Phys. 07 (2013) 025.

[62] J. L. Sievers, R. A. Hlozek, M. R. Nolta, V. Acquaviva, G. E. Addison, P. A. R. Ade, P. Aguirre, M. Amiri, J. W. Appel, L. F. Barrientos et al., J. Cosmol. Astropart. Phys. 10 (2013) 060.

[63] P. A. R. Ade, N. Aghanim, C. Armitage-Caplan, M. Arnaud, M. Ashdown, F. Atrio-Barandela, J. Aumont, C. Baccigalupi, A. J. Banday et al. (Planck Collaboration), Astron. Astrophys. **571**, A21 (2014).

[64] E. M. George, C. L. Reichardt, K. A. Aird, B. A. Benson, L. E. Bleem, J. E. Carlstrom, C. L. Chang, H. M. Cho, T. M. Crawford, A. T. Crites et al., Astrophys. J. **799**, 177 (2015).

[65] N. Aghanim, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, J. G. Bartlett, N. Bartolo et al. (Planck Collaboration), Astron. Astrophys. **594**, A22 (2016).

[66] D. Han, N. Sehgal, and F. Villaescusa-Navarro, Phys. Rev. D **104**, 123521 (2021).

[67] P. A. R. Ade, N. Aghanim, C. Armitage-Caplan, M. Arnaud, M. Ashdown, F. Atrio-Barandela, J. Aumont, C. Baccigalupi, A. J. Banday et al. (Planck Collaboration), Astron. Astrophys. **571**, A18 (2014).

[68] J. C. Hill and D. N. Spergel, J. Cosmol. Astropart. Phys. 02 (2014) 030.

[69] A. Lewis, A. Challinor, and A. Lasenby, Astrophys. J. **538**, 473 (2000).

[70] M. Kamionkowski, A. Kosowsky, and A. Stebbins, Phys. Rev. D **55**, 7368 (1997).

[71] N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo et al. (Planck Collaboration), Astron. Astrophys. **641**, A6 (2020).

[72] CMB-HD Survey and Instrument Overview, https://cmb-hd.org/survey-and-instrument-overview (accessed: Nov 2nd 2021).

[73] C. Prasad and N. Sehgal (to be published).

[74] S. Raghunathan, N. Whitehorn, M. A. Alvarez, H. Aung, N. Battaglia, G. P. Holder, D. Nagai, E. Pierpaoli, C. L. Reichardt, and J. D. Vieira, Astrophys. J. **926**, 172 (2022).

[75] K. M. Smith, M. S. Madhavacheril, M. Münchmeyer, S. Ferraro, U. Giri, and M. C. Johnson, arXiv:1810.13423.

[76] J. Cayuso, R. Bloch, S. C. Hotinli, M. C. Johnson, and F. McCarthy, arXiv:2111.11526.

[77] Ž. Ivezić, S. M. Kahn, J. A. Tyson, B. Abel, E. Acosta, R. Allsman, D. Alonso, Y. AlSayyad, S. F. Anderson, J. Andrew *et al.*, Astrophys. J. **873,** 111 (2019).

[78] J. Caldeira, W. L. K. Wu, B. Nord, C. Avestruz, S. Trivedi, and K. T. Story, Astron. Comput. **28,** 100307 (2019).

[79] M. A. Petroff, G. E. Addison, C. L. Bennett, and J. L. Weiland, Astrophys. J. **903,** 104 (2020).

[80] E. Guzman and J. Meyers, Phys. Rev. D **104,** 043529 (2021).

[81] E. Guzman and J. Meyers, J. Cosmol. Astropart. Phys. 01 (2022) 030.

[82] Z. Lin, N. Huang, C. Avestruz, W. L. K. Wu, S. Trivedi, J. Caldeira, and B. Nord, Mon. Not. R. Astron. Soc. **507,** 4149 (2021).

[83] P. Villanueva-Domingo and F. Villaescusa-Navarro, Astrophys. J. **907,** 44 (2021).

[84] H. Cai, M. S. Madhavacheril, J. C. Hill, and A. Kosowsky, Phys. Rev. D **105,** 043516 (2022).

[85] W. Hu, Phys. Rev. D **62,** 043007 (2000).

[86] W. Hu and T. Okamoto, Astrophys. J. **574,** 566 (2002).

[87] B. D. Sherwin, A. van Engelen, N. Sehgal, M. Madhavacheril, G. E. Addison, S. Aiola, R. Allison, N. Battaglia, D. T. Becker, J. A. Beall *et al.*, Phys. Rev. D **95,** 123529 (2017).

[88] N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo *et al.* (Planck Collaboration), Astron. Astrophys. **641,** A8 (2020).

[89] W. L. K. Wu, L. M. Mocanu, P. A. R. Ade, A. J. Anderson, J. E. Austermann, J. S. Avva, J. A. Beall, A. N. Bender, B. A. Benson, F. Bianchini *et al.*, Astrophys. J. **884,** 70 (2019).

[90] W. Hu, S. DeDeo, and C. Vale, New J. Phys. **9,** 441 (2007).

[91] S. Dodelson and G. D. Starkman, arXiv:astro-ph/0305467.

[92] G. Holder and A. Kosowsky, Astrophys. J. **616,** 8 (2004).

[93] A. Lewis and L. King, Phys. Rev. D **73,** 063006 (2006).

[94] M. Madhavacheril, N. Sehgal, R. Allison, N. Battaglia, J. R. Bond, E. Calabrese, J. Caligiuri, K. Coughlin, D. Crichton, R. Datta *et al.*, Phys. Rev. Lett. **114,** 151302 (2015).

[95] M. S. Madhavacheril, J. C. Hill, S. Næss, G. E. Addison, S. Aiola, T. Baildon, N. Battaglia, R. Bean, J. R. Bond, E. Calabrese *et al.*, Phys. Rev. D **102,** 023534 (2020).

[96] B. Hadzhiyska, B. D. Sherwin, M. Madhavacheril, and S. Ferraro, Phys. Rev. D **100,** 023547 (2019).

[97] B. Horowitz, S. Ferraro, and B. D. Sherwin, Mon. Not. R. Astron. Soc. **485,** 3919 (2019).

[98] T. Namikawa, D. Hanson, and R. Takahashi, Mon. Not. R. Astron. Soc. **431,** 609 (2013).

[99] M. Kesden, A. Cooray, and M. Kamionkowski, Phys. Rev. D **67,** 123507 (2003).

[100] K. T. Story, D. Hanson, P. A. R. Ade, K. A. Aird, J. E. Austermann, J. A. Beall, A. N. Bender, B. A. Benson, L. E. Bleem, J. E. Carlstrom *et al.*, Astrophys. J. **810,** 50 (2015).

[101] D. J. E. Marsh, arXiv:1605.05973.

[102] M. Aravena, J. Austermann, K. Basu, N. Battaglia, B. Beringue, F. Bertoldi, J. R. Bond, P. Breysse, R. Bustos, S. Chapman *et al.*, Bull. Am. Astron. Soc. **51,** 213 (2019).

[103] E. Schaan, S. Ferraro, M. Vargas-Magaña, K. M. Smith, S. Ho, S. Aiola, N. Battaglia, J. R. Bond, F. De Bernardis, E. Calabrese *et al.*, Phys. Rev. D **93,** 082002 (2016).

[104] E. Schaan, S. Ferraro, S. Amodeo, N. Battaglia, S. Aiola, J. E. Austermann, J. A. Beall, R. Bean, D. T. Becker, R. J. Bond *et al.*, Phys. Rev. D **103,** 063513 (2021).

[105] L. Knox, Phys. Rev. D **52,** 4307 (1995).