

## Avoiding baryonic feedback effects on neutrino mass measurements from CMB lensing

Fiona McCarthy,<sup>1,2</sup> Simon Foreman,<sup>1,3</sup> and Alexander van Engelen<sup>4</sup>

<sup>1</sup>*Perimeter Institute for Theoretical Physics, Waterloo, Ontario N2L 2Y5, Canada*

<sup>2</sup>*Department of Physics and Astronomy, University of Waterloo, Waterloo, Ontario N2L 3G1, Canada*

<sup>3</sup>*Dominion Radio Astrophysical Observatory, Herzberg Astronomy and Astrophysics Research Centre, National Research Council Canada, P.O. Box 248, Penticton, British Columbia V2A 6J9, Canada*

<sup>4</sup>*School of Earth and Space Exploration, Arizona State University, Tempe, Arizona 85287, USA*



(Received 2 December 2020; accepted 30 April 2021; published 28 May 2021)

A measurement of the sum of neutrino masses is one of the main applications of upcoming measurements of gravitational lensing of the cosmic microwave background (CMB). This measurement can be confounded by modelling uncertainties related to so-called “baryonic effects” on the clustering of matter, arising from gas dynamics, star formation, and feedback from active galactic nuclei and supernovae. In particular, a wrong assumption about the form of baryonic effects on CMB lensing can bias a neutrino mass measurement by a significant fraction of the statistical uncertainty. In this paper, we investigate three methods for mitigating this bias: (1) restricting the use of small-scale CMB lensing information when constraining neutrino mass; (2) using an external tracer to remove the low-redshift contribution to a CMB lensing map; and (3) marginalizing over a parametric model for baryonic effects on large-scale structure. We test these methods using Fisher matrix forecasts for experiments resembling the Simons Observatory and CMB-S4, using a variety of recent hydrodynamical simulations to represent the range of possible baryonic effects, and using cosmic shear measured by the Rubin Observatory’s LSST as the tracer in method (2). We find that a combination of (1) and (2), or (3) on its own, will be effective in reducing the bias induced by baryonic effects on a neutrino mass measurement to a negligible level, without a significant increase in the associated statistical uncertainty.

DOI: [10.1103/PhysRevD.103.103538](https://doi.org/10.1103/PhysRevD.103.103538)

### I. INTRODUCTION

With improving measurements of the cosmic microwave background (CMB), a cosmological measurement of the sum of the neutrino masses is envisioned within the next decade. Within the standard model of particle physics, the three neutrinos are massless particles; thus, the first measurements of neutrino flavor oscillations [1,2], a process that only occurs if there exist mass differences between the species, were key developments in the search for beyond-standard-model physics. Neutrino oscillation experiments are sensitive to the difference in the squares of the masses of neutrinos,  $\Delta m_{ij}^2 \equiv m_i^2 - m_j^2$ ; however, a cosmological neutrino detection will be sensitive to the sum of the neutrino masses  $M_\nu \equiv \sum_{i=1}^3 m_i$ , and thus will be important in setting the overall scale of the neutrino masses. The current lower limit on  $M_\nu$  (from neutrino oscillation experiments) is  $M_\nu \gtrsim 60$  meV [3]. Until now, cosmological experiments have only placed upper limits on  $M_\nu$ ; the best is that of the Planck survey [4], which gives  $M_\nu < 120$  meV.

Massive neutrinos have a well-understood effect on the matter power spectrum  $P_m(k, z)$ . After becoming

nonrelativistic when their temperature  $T_\nu$  was comparable to their mass, they started contributing to  $P_m(k, z)$ ; however, due to their small masses, they do not cluster on small scales, instead free streaming, leading to a suppression of power on small scales. See [5,6] for reviews on the cosmological effects of neutrinos.

Planck’s upper bound on the neutrino mass was obtained from a joint analysis of the CMB temperature and polarization maps, CMB lensing maps, and baryonic acoustic oscillation (BAO) measurements. As the neutrinos were still relativistic at the time of recombination when the CMB was released, the majority of a CMB survey’s constraining power on  $M_\nu$  comes from the CMB lensing information, which is sensitive to large-scale structure at all redshifts. In coming years, experiments such as the Simons Observatory (SO) [7], SPT-3G [8], and, further in the future, CMB-S4 [9], will make better measurements of the CMB lensing power spectrum, and have been forecast to measure the neutrino mass to between 20 and 30 meV [7,9,10].

To be able to reach this level of constraint, the CMB lensing power spectrum must be well understood theoretically. In particular, the lensing power spectrum is a projection of the matter power spectrum  $P_m(k, z)$  over

all redshifts; however, there are certain effects that currently limit our understanding of  $P_m(k, z)$ , in particular effects due to baryonic processes (such as gas cooling and feedback from supernovae and active galactic nuclei) in the Universe. Most predictions of  $P_m(k, z)$  only account for gravitational forces, neglecting the complex baryonic interactions that we know exist; our current best method for understanding  $P_m(k, z)$  including baryonic physics is performing large hydrodynamical simulations. Measurements of  $P_m(k, z)$  from different simulations differ due to different numerical schemes and phenomenological implementations of baryonic processes that cannot be directly simulated at a given resolution. While the “true” impact of baryonic effects on the matter power spectrum is not known, a general conclusion is that baryons contribute to a suppression of power on small scales [e.g.,  $0.1h \text{ Mpc}^{-1} \lesssim k \lesssim \mathcal{O}(\text{few} \times 10)h \text{ Mpc}^{-1}$  at  $z = 0$ ].

Uncertainty due to baryonic feedback has been extensively studied in the context of cosmic shear surveys [11–18], as the scales affected by feedback directly correspond to the scales that current shear surveys are most sensitive to. However, Ref. [16] and recently Ref. [19] found that uncertainty from baryonic effects can also be important in the search for neutrino masses from CMB lensing, in spite of the higher redshifts and larger length scales involved. If we are to trust a measurement of  $M_\nu$  from CMB lensing, it will be important to have an inference which is robust to these baryonic effects. Furthermore, the sensitivity of CMB lensing to baryonic effects implies that we could learn about the latter from observations of the former, an avenue explored for cosmic shear in Refs. [20–23].

Reference [19] explored the lensing power spectrum suppression and associated bias on neutrino mass from a suite of recent hydrodynamic simulations, finding that the range in possible biases is non-negligible compared with expected statistical uncertainties. They found a significant scatter between different simulations, comparable to the statistical uncertainty in the measurement. In this paper, we consider various methods of mitigating this bias on the inference of  $M_\nu$  from CMB lensing surveys similar to SO and CMB-S4. We describe various techniques to remove the sensitivity to the relevant baryonic processes and test, using Fisher forecasts, how these techniques will reduce the bias for the series of simulations examined by Ref. [19]. The first mitigation method we consider is a simple scale cut, where the smallest scales of the CMB lensing convergence (which are most sensitive to baryonic effects) are removed from the analysis. Second, we consider a method of using external cosmic shear measurements to estimate and subtract the low- $z$  contribution to the CMB lensing potential—as baryonic effects are relatively late-Universe phenomena, we expect the high- $z$  portion of the CMB lensing potential to be less sensitive to them, while still retaining sensitivity to  $M_\nu$ . Finally, we ask whether

marginalizing over the parameters of a general model for baryonic effects will reduce the bias while preserving the  $M_\nu$  constraints.

We find that imposing a scale cut of  $L_{\text{max}} \sim 1000$  on the lensing multipoles used for constraining  $M_\nu$  can reduce the bias from baryonic effects by up to a factor of 2, with more aggressive cuts significantly increasing the statistical uncertainty. On the other hand, combining this scale cut with subtraction of a low- $z$  tracer, or marginalizing over a baryonic model (with or without a scale cut) will be much more effective in eliminating the bias, reducing it by at least a factor of 5 in the first case and 10 in the second case for all simulations we consider. The maximum residual bias associated with these simulations is  $\sim 3 \text{ meV}$  in either case, well below the level that would interfere with a high-significance detection of the minimum allowed neutrino mass sum.

The paper is organized as follows. In Sec. II we discuss the CMB lensing power spectrum and review the range of possible baryonic effects as represented by current hydrodynamical simulations. In Sec. III we present our Fisher forecast formalism for calculating the forecast constraint and biases on the inference of  $M_\nu$  from these simulations. In Sec. IV we discuss the effects of a small angular-scale cutoff in  $C_L^{\kappa\kappa}$ . In Sec. V we discuss the effect of “subtracting” a low- $z$  tracer to isolate the high- $z$  contribution to the CMB lensing map. In Sec. VI we discuss the effect of marginalizing over parameters that describe the baryonic effects on the matter power spectrum. We discuss our results in Sec. VII.

## II. CMB LENSING AND BARYONS

The CMB photons we detect have been gravitationally lensed by any matter they encounter along the paths they have travelled since their “release” during recombination, at  $z \sim 1100$ . Structures at any redshift after recombination can act as lenses, making CMB lensing a powerful probe of the evolution of the matter content of the Universe. We mainly quantify this information via the convergence power spectrum  $C_L^{\kappa\kappa}$ , which is a line-of-sight integral over the matter power spectrum  $P_m(k, z)$  (see e.g., [24]),

$$C_L^{\kappa\kappa} = \int_0^{\chi_{\text{CMB}}} d\chi \frac{W_{\text{CMB}}^\kappa(\chi)^2}{\chi^2} P_m\left(k = \frac{L + 1/2}{\chi}, z\right), \quad (1)$$

where  $W_{\text{CMB}}^\kappa(\chi)$  is the CMB lensing efficiency kernel

$$W_{\text{CMB}}^\kappa(\chi) = \frac{3}{2} \Omega_m \left(\frac{H_0}{c}\right)^2 \frac{\chi}{a(\chi)} \frac{\chi_{\text{CMB}} - \chi}{\chi_{\text{CMB}}}, \quad (2)$$

with  $H_0$  the Hubble constant today,  $c$  the speed of light,  $\Omega_m$  the density of matter today,  $\chi_{\text{CMB}}$  the comoving distance to the surface of last scattering (at which the CMB was released), and  $a(\chi) = \frac{1}{1+z(\chi)}$  and the scale factor at

comoving distance  $\chi$ . Equation (1) assumes the Limber approximation [25,26], which is valid in the small-scale, flat-sky limit; and the Born approximation, where the integral is taken over the photon's undeflected path, valid in the small-deflection limit [27,28].

To make an accurate inference of  $M_\nu$  from a CMB lensing survey, we need to trust our theoretical model of the lensing convergence; i.e., we need to understand every component of Eq. (1). The cosmological ingredients that enter CMB lensing kernel  $W^k(\chi)$  are well understood. On the other hand,  $P_m(k, z)$  is most commonly computed from linear gravitational perturbation theory on large scales, supplemented on small scales by nonlinear extensions of gravitational perturbation theory, phenomenological models,  $N$ -body simulations, or emulators. Generally, these only account for gravitational interactions between the matter; i.e., they treat all matter as “dark.”

However, about 15% of matter is not dark but baryonic, and has complex interactions with itself and with light. These interactions effect changes to how matter clusters on  $\sim$ Mpc and smaller scales: as examples, gas cooling and AGN feedback cause matter to condense and expand respectively (e.g., [29]). Our models currently lack a first-principles calculation of the power spectrum  $P_m(k, z)$  incorporating these interactions, and they are typically neglected when considering CMB lensing surveys. However, some of the baryonic effects on  $P_m(k, z)$ —particularly the suppression of power on small scales—mimic the effects of massive neutrinos, and neglecting these in the theoretical modeling of  $C_L^{\kappa\kappa}$  can lead to significant biases on the neutrino mass inference [16,19].

In Fig. 1, we illustrate the suppression of the CMB lensing power spectrum by baryonic effects, as computed in Ref. [19]<sup>1</sup> for a selection of recent hydrodynamical simulations (see Sec. III C). At large scales, the power spectra coincide, while the baryonic suppression becomes relevant at  $L \gtrsim 1000$ . We also show the effect of a nonzero neutrino mass on the lensing power spectrum, by plotting the ratio of the fiducial  $M_\nu = 60$  meV power spectrum to one where  $M_\nu = 0$  meV (with all other cosmological parameters unchanged). In this case, we see a power suppression with much milder scale dependence than for baryonic effects. This difference indicates that it may be possible to disentangle the two types of suppression, motivating the methods we consider in this work.

### III. NEUTRINO MASS: CONSTRAINTS AND BIAS

#### A. Forecasting the $1\sigma$ constraints

We consider an analysis where  $M_\nu$  is allowed to vary along with other cosmological parameters; as such we consider a parameter vector




FIG. 1. The effect of baryons on the CMB lensing power spectrum  $C_L^{\kappa\kappa}|_{\text{bary}}$  in a selection of large hydrodynamical simulations (see Sec. III C), as computed in Ref. [19], shown as a ratio with the spectrum for dark matter only,  $C_L^{\kappa\kappa}|_{\text{DMO}}$ . The effect of nonzero  $M_\nu$  is also shown. The different scale dependences of baryonic and neutrino-mass effects indicate that it may be possible to distinguish between the two in lensing measurements.

$$\vec{\theta} = (h, \Omega_b h^2, \Omega_c h^2, \tau, n_s, A_s, M_\nu) \quad (3)$$

with fiducial values  $\{h = 0.675, \Omega_b h^2 = 0.0222, \Omega_c h^2 = 0.1197, \tau = 0.06, n_s = 0.9655, A_s = 2.2 \times 10^{-9}, M_\nu = 0.06 \text{ eV}\}$  corresponding to the best-fit parameters of the Planck analysis [4] for the first six parameters, namely, the Hubble parameter in units of  $100 \text{ km s}^{-1} \text{ Mpc}^{-1}$ , the physical baryon density, the physical cold dark matter density, the optical depth to recombination, the primordial scalar fluctuation slope and amplitude (with a pivot scale of  $0.05 \text{ Mpc}^{-1}$ ). We take a fiducial value of 60 meV (the minimum allowed value) for the sum of the neutrino mass  $M_\nu$ . We compute the fiducial matter power spectrum with CAMB [31], with the nonlinearities and treatment of neutrinos given by the extended halo model from Ref. [32].

As the cosmological parameters will also be constrained from the primary CMB, we include information from the primary CMB temperature and polarization as measured by the experiment we are forecasting for (see Sec. III D), as well as a prior from BAO measurements from DESI [33], which improves the analysis by breaking the geometric degeneracy in the CMB.

To compute the information from the CMB and CMB lensing, we use the Fisher formalism, in which the Fisher matrix  $F$  approximates the inverse covariance matrix of the parameters, with the diagonals of  $F^{-1}$  giving the squares of the expected  $1\sigma$  uncertainties on each parameter (with all other parameters marginalized over). We calculate  $F$  according to

<sup>1</sup>These computations are available from Ref. [30].

$$F_{ij}^{\text{CMB}} = \sum_{\ell} \frac{2\ell + 1}{2} f_{\text{sky}} \text{Tr} \left[ \frac{\partial C_{\ell}^{\text{CMB}}}{\partial \theta^i} (C_{\ell}^{\text{CMB}})^{-1} \times \frac{\partial C_{\ell}^{\text{CMB}}}{\partial \theta^j} (C_{\ell}^{\text{CMB}})^{-1} \right]. \quad (4)$$

In Eq. (4),  $f_{\text{sky}}$  is the fraction of sky area which the surveys cover and  $C_{\ell}^{\text{CMB}}$  is the covariance matrix of the CMB:

$$C_{\ell}^{\text{CMB}} = \begin{pmatrix} C_{\ell}^{TT} & C_{\ell}^{TE} & C_{\ell}^{T\kappa} \\ C_{\ell}^{TE} & C_{\ell}^{EE} & C_{\ell}^{E\kappa} \\ C_{\ell}^{T\kappa} & C_{\ell}^{E\kappa} & C_{\ell}^{\kappa\kappa} \end{pmatrix}, \quad (5)$$

where  $C_{\ell}^{TT}$  is the power spectrum of the observed temperature anisotropies (including noise);  $C_{\ell}^{EE}$  is the power spectrum of the observed  $E$ -mode polarization anisotropies (also including noise); and  $C_{\ell}^{TE}$  is their cross power spectrum.  $C_{\ell}^{\kappa\kappa}$  includes the reconstruction noise for CMB lensing. Although the CMB we measure is lensed, we use the unlensed primary CMB power spectra  $C_{\ell}^{TT}$ ,  $C_{\ell}^{EE}$ , and  $C_{\ell}^{TE}$  to avoid double counting of the lensing information. A proper treatment including lensed CMB power spectra would involve including the covariances between the CMB power spectra induced by lensing, and also the covariances between the lensing convergence and the CMB power spectra [34,35]; neglecting the extra information on  $M_{\nu}$  that comes from the lensed CMB power spectra makes our calculation conservative.  $C_{\ell}^{T\kappa}$  and  $C_{\ell}^{E\kappa}$ , the cross power spectra of CMB lensing with CMB temperature and  $E$ -mode polarization, respectively, are nonzero only on very large scales due to correlations induced by the late-Universe effects on the CMB such as the integrated Sachs-Wolfe (ISW) effect [36] and polarization generated after reionization [37], but can be neglected in our analysis as we restrict to multipoles  $\ell > 300$  for the primary CMB.

We also include a prior on  $\tau$ , the optical depth to reionization, which will be an important limiting factor in the inference of  $M_{\nu}$  from lensing surveys [10]. We consider two different scenarios: the Planck design sensitivity  $\sigma_{\text{prior}}(\tau) = 0.006$  (equal to the value achieved by the analysis of Planck data in Ref. [38]) and the cosmic variance limit  $\sigma_{\text{prior}}(\tau) = 0.002$  [9]. We include this as a Gaussian prior with width  $\sigma_{\text{prior}}(\tau)$ . The final Fisher matrix we use for forecasting is

$$F = F^{\text{CMB}} + C_{\text{prior}}^{-1} \quad (6)$$

where  $C_{\text{prior}}$  is the sum of the BAO prior and the  $\tau$  prior. Note that  $F^{\text{CMB}}$  contains both primary CMB and lensing information, as it is calculated from the covariance matrix in Eq. (5); however, as the cross power spectra between the primary CMB and the CMB lensing potential are set to

zero, this can be separated as a sum of an inverse prior from the primary CMB and a Fisher matrix due to lensing alone.

Within this setup, the lower bound on the marginalized constraint on parameter  $i$  is

$$\sigma_i = \sqrt{(F^{-1})_{ii}}. \quad (7)$$

## B. Calculating the baryonic bias

Predictions for the CMB lensing power spectrum are typically computed with the dark-matter-only (DMO)<sup>2</sup> nonlinear matter power spectrum  $P_{\text{DMO}}(k, z)$ . However, since the true power spectrum includes (unknown) baryonic effects, the deviation from the DMO prediction that these effects induce might mimic the neutrino mass signal and result in an incorrect (“biased”) inference of the mass. If we can compute the power spectrum incorporating a given model of baryonic effects [ $P_{\text{bary}}(k, z)$ ], we can calculate the bias that would be induced in the inference of the parameter  $\theta^i$  by (e.g., [16])

$$B_i = F^{-1} \sum_{\ell} \frac{2\ell + 1}{2} f_{\text{sky}} \text{Tr} \left[ \frac{\partial C_{\ell}^{\text{CMB}}}{\partial \theta^i} (C_{\ell}^{\text{CMB}})^{-1} \times \Delta C_{\ell} (C_{\ell}^{\text{CMB}})^{-1} \right], \quad (8)$$

where  $\Delta C_{\ell}$  is the change in the covariance matrix due to baryonic effects

$$\Delta C_{\ell} \equiv C_{\ell}|_{\text{bary}} - C_{\ell}|_{\text{DMO}}, \quad (9)$$

and where, as in the previous subsection, the covariance matrix  $C_{\ell}^{\text{CMB}}$  includes the instrumental noise for the given survey. Note that of all the power spectra in the covariance matrix (5), only  $C_{\ell}^{\kappa\kappa}$  is affected by the baryons, and  $\Delta C_{\ell}^{XY} = 0$  for  $XY \neq \kappa\kappa$  (the  $\kappa T$  and  $\kappa E$  correlations are too small to be relevant). Thus, Eq. (8) simplifies to

$$B_i = F^{-1} \sum_{\ell} \frac{2\ell + 1}{2} f_{\text{sky}} \frac{\partial C_{\ell}^{\kappa\kappa}}{\partial \theta^i} \frac{1}{(C_{\ell}^{\kappa\kappa})^2} \Delta C_{\ell}^{\kappa\kappa}. \quad (10)$$

We use the forecasting code from Ref. [39]<sup>3</sup> to compute the Fisher matrices and biases in our forecasts. Note that the bias depends on the experiment, because both the Fisher

<sup>2</sup>“Dark-matter-only” computations could perhaps be more accurately described as “gravity only,” since these computations do not neglect the baryonic contribution to the Universe’s matter content but instead treat baryonic identically to dark matter, with only gravitational forces at play. However, “dark matter only” is the term most commonly seen in the literature, so we also adopt it in this work.

<sup>3</sup>See Ref. [40].



matrix and  $(C_{\ell}^{\kappa\kappa})^{-2}$  factors in Eq. (10) depend on the assumed beam and noise.

To obtain a range of possible forms for  $P_{\text{bary}}$ , and therefore  $\Delta C_{\ell}^{\kappa\kappa}$ , we turn to hydrodynamical simulations, as described in the next subsection.

### C. Simulations

We use  $C_L^{\kappa\kappa}$  computations from Ref. [19], which considers seven baryonic scenarios from four different families of large hydrodynamical simulations (see Ref. [19] for further descriptions):

- (i) The ‘‘AGN’’ member of the OWLS simulation suite [41–43].
- (ii) The base BAHAMAS simulation [43–45], along with the ‘‘Low-AGN’’ and ‘‘High-AGN’’ versions that, respectively, contain weaker and stronger AGN feedback than the base simulation.
- (iii) The ‘‘AGN’’ member of the Horizon simulation suite [46–48].
- (iv) The TNG100 and TNG300 runs of the IllustrisTNG simulations [49–54].

The matter power spectrum  $\hat{P}(k, z)$  is measured from the simulation outputs at several different redshifts, both from DMO runs (which treat baryons and dark matter identically) and from runs that include baryonic processes along with gravity. The measured power spectra have considerable uncertainty due to sample variance arising from the finite number of modes within each simulated volume, but the majority of this sample variance arises from randomness in the initial conditions that manifests primarily at large scales. Each pair of DMO and full-hydro runs begins with the same initial conditions (i.e., amplitudes and phases of modes at the initial time), and therefore the sample variance errors mostly cancel<sup>4</sup> in the ratio

$$\hat{R}(k, z) \equiv \frac{\hat{P}_{\text{bary}}(k, z)}{\hat{P}_{\text{DMO}}(k, z)}. \quad (11)$$

The corresponding CMB lensing power spectrum  $C_{\ell}^{\kappa\kappa}|_{\text{bary}}$  can then be computed by using

$$P_{\text{bary}}(k, z) = P_{\text{fid}}(k, z)\hat{R}(k, z) \quad (12)$$

<sup>4</sup>The ratio  $\hat{R}(k, z)$  will itself have some sample variance, because it is dominated by baryonic effects on the highest-mass halos within a given simulation volume, and the set of such halos will depend on the initial conditions. Reference [55] quantified the sample variance in  $\hat{R}(k, z)$  for a subset of the simulations considered in this work, finding it to be at the few-percent level for  $k \lesssim 20h \text{ Mpc}^{-1}$  (Refs. [43,48] reached similar conclusions.). This is acceptable for our work, which is focused on the range of  $\hat{R}(k, z)$  between different simulations rather than the absolute precision of any one simulation.

in Eq. (1), where  $P_{\text{fid}}(k, z)$  is the fiducial DMO prediction for the matter power spectrum. Note that the different simulations have been run with different cosmological models, while we compute  $P_{\text{fid}}$  using a single cosmology in our forecasts. References [43,56,57] have found that  $\hat{R}(k, z)$  has only a weak dependence on background cosmology, so Eq. (12) is sufficient for our forecasts, while for work requiring percent-level accuracy, the cosmology dependence of  $\hat{R}(k, z)$  should be carefully accounted for.<sup>5</sup>

### D. Experimental configurations

Several CMB experiments are planned or being built that will begin observations this decade and that are aimed at measuring CMB fluctuations on small scales, for gravitational lensing reconstruction and other secondary anisotropies. Here, we consider an experiment similar to the SO [7], due to begin taking data in the first half of the 2020s. The large aperture telescope for this experiment will have a 6 m diameter and will observe large fractions of the sky at high angular resolution in six frequency channels. We also consider an experiment like CMB-S4 [9,59], which will have comparable angular resolution and frequency coverage but higher sensitivity; it is expected to begin taking data on a later timeline than SO. We include Gaussian instrumental white noise on the CMB power spectra:

$$N_{\ell} = N_T e^{\frac{\ell(\ell+1)\Theta_{\text{FWHM}}^2}{8 \ln 2}}, \quad (13)$$

where  $N_T$  is the noise variance and  $\Theta_{\text{FWHM}}$  is the beam size of the experiment. For both experiments, we use  $\Theta_{\text{FWHM}} = 1.4'$ ; for SO, we use  $N_T$  corresponding to a map noise level of  $6 \mu\text{K-arcmin}$  and for S4 we use  $N_T$  corresponding to  $1 \mu\text{K-arcmin}$ .

For  $C_L^{\kappa\kappa}$ , we include reconstruction noise  $N_L^{\kappa\kappa}$  corresponding to the minimum-variance reconstruction of Ref. [60] relevant to the experiment we are considering (SO or S4). We include multipoles  $90 \leq L \leq 3100$  in the lensing power spectra, with the upper limit chosen based on where the statistical sensitivity drops off, and the lower limit having negligible impact on the results. For the primary CMB, we include multipoles  $300 \leq \ell \leq 3000$  for  $C_{\ell}^{TT}$  and  $300 \leq \ell \leq 5000$  for  $C_{\ell}^{EE}$ , with the upper limits based on where uncleaned foregrounds are expected to become significant in the lensing reconstruction [61,62]. We assume a sky fraction of  $f_{\text{sky}} = 0.4$  and full overlap between all the fields we consider. For the  $EB$ -based reconstruction, which dominates the information at S4 noise, we include iterated delensing [63].

<sup>5</sup>An example is precise comparisons of  $P_{\text{bary}}$  to predictions from perturbation theory [58]. We thank Matthew Lewandowski for discussions on this point.

#### IV. STRATEGY 1: SMALL ANGULAR-SCALE CUTOFF

The forecasts in Ref. [19] considered all scales over which future CMB surveys will have appreciable sensitivity to the lensing power spectrum—i.e., a summation over multipoles  $90 < L < 3100$ . However, baryonic effects are concentrated at a different (though not disjoint) range of scales than the neutrino mass constraint (see Fig. 1).

With this in mind, the first mitigation strategy we implement is a simple  $L_{\max}$  cutoff, where  $L_{\max}$  is the maximum lensing multipole included in the analysis. In Fig. 2, we show the behavior of the  $1\sigma$  constraint on  $M_\nu$  as well as the biases from different models of baryonic effects as we introduce this cutoff. It is clear that for all of the experimental setups, the constraints saturate at around  $L_{\max} \sim 1000$  and there is no benefit to including multipoles  $L \gtrsim 1000$ ; this happens because the suppression of  $C_L^{\kappa\kappa}$  is roughly constant for  $L \gtrsim 1000$ , while the experimental error bars on  $C_L^{\kappa\kappa}$  increase with  $L$  over the same range (compare Fig. 1 of this work with Fig. 1 of Ref. [19]). Meanwhile, we see that including higher multipoles does indeed increase the bias, and imposing  $L_{\max} \sim 1000$  can reduce the bias by a factor of  $\sim 2$  in some cases. However, in particular for the most advanced experimental configuration, the biases can still be of the same order of magnitude as the expected constraint, and so further mitigation methods will be needed to reduce the bias to an acceptable level.

#### V. STRATEGY 2: SUBTRACTION OF EXTERNAL TRACERS

##### A. Isolating the low- $z$ contribution to the CMB lensing potential

Baryonic phenomena begin to imprint themselves on structure formation at a much later time in cosmological

history than neutrino mass effects. As the CMB lensing kernel is an integral over all redshifts, we receive (weighted) information from all of cosmological history since recombination. However, if we could “subtract” the low- $z$  contribution to the lensing map to isolate the high- $z$  effects, we could potentially remove most of the bias while still being sensitive to  $M_\nu$ .

To illustrate the ideal outcome of such a procedure, we define a high- $z$  CMB lensing field by

$$C_L^{\kappa_h \kappa_h} = \int_{\chi_{\min}}^{\chi_{\text{CMB}}} d\chi \frac{W^\kappa(\chi)^2}{\chi^2} P_m \left( k = \frac{L + 1/2}{\chi}, z \right), \quad (14)$$

where  $\chi_{\min}$  is some lower bound of the integration;  $\chi_{\min} = 0$  corresponds to the standard CMB lensing scenario. We forecast the  $1\sigma$  errors and baryonic biases on  $M_\nu$  as in Sec. III but replacing  $C_L^{\kappa\kappa}$  with  $C_L^{\kappa_h \kappa_h}$ , and we consider their dependences on the lower limit of integration  $\chi_{\min}$ .

We show in Fig. 3 the behavior of the constraints and the biases plotted against  $z_{\min} = z(\chi_{\min})$ . Here we see explicitly that the biases are introduced in the late Universe at around  $z \lesssim 2$ , while the constraints on  $M_\nu$  come from a much larger redshift range. It is clear that if we could isolate the portion of the lensing map that is sourced at  $z \gtrsim 2$ , then we could remove a significant portion of the baryonic bias on  $M_\nu$  without sacrificing much constraining power on  $M_\nu$ .

Of course, we do not have direct access to measurements of the field  $\kappa_h$  and so the configuration in Fig. 3 is simply a toy model for illustrative purposes. However, through cross-correlation with a low- $z$  external tracer  $\hat{X}$  (such as a cosmic shear map), we could remove the low- $z$  portion of the CMB lensing field by defining a new field




FIG. 2. Demonstrating mitigation strategy 1: the calculated biases on  $M_\nu$  plotted against the maximum CMB lensing multipole  $L_{\max}$  included in the forecast for the different simulations. In each case the  $1\sigma$  constraint is shown as a dotted line. Note that the TNG100 and TNG300 lines are almost identical.




FIG. 3. The biases (solid lines) and the forecast  $1\sigma$  constraints on  $M_\nu$ , plotted against the minimum  $z$  ( $z_{\min}$ ) used to integrate the CMB lensing kernel, assuming that perfect removal of the  $z < z_{\min}$  contribution is possible. We see that the baryonic bias starts to become relevant at  $z_{\min} \sim 2$ , and that the bias is larger for S4 than for SO due to its higher sensitivity at small scales. However, constraining information for  $M_\nu$  comes from higher  $z$ . Note that the  $x$  axis is scaled logarithmically in  $(1+z)$ .

$$\hat{\kappa}_{\text{sub}}^{\text{sub}}(\mathbf{L}) = \hat{\kappa}(\mathbf{L}) - \frac{C_L^{\chi\kappa} + N_L^{\chi\kappa}}{C_L^{\chi\chi} + N_L^{\chi\chi}} \hat{X}(\mathbf{L}), \quad (15)$$

where  $\hat{\kappa}(\mathbf{L})$  is the original CMB lensing convergence, and in this subsection we use hats to denote quantities that include noise (i.e.,  $C_L^{\hat{A}\hat{B}} = C_L^{AB} + N_L^{AB}$ ). In Eq. (15), we weight  $\hat{X}$  with a matched filter designed to extract the portion of  $X$  correlated with  $\kappa$ , assuming that this correlation is dominated by low redshifts. Note that, in our implementation, we will assume that the filter is computed using theoretical expressions with the fixed fiducial cosmological parameters, such that it is not varied in our Fisher calculations.

To see how Eq. (15) accomplishes our goal, observe that if the weights perfectly match the true statistics of  $\hat{X}$  and  $\hat{\kappa}$ , the power spectrum of  $\hat{\kappa}_{\text{sub}}$  reduces to

$$C_L^{\hat{\kappa}_{\text{sub}}\hat{\kappa}_{\text{sub}}} = C_L^{\kappa\kappa} + N_L^{\kappa\kappa} - \frac{(C_L^{\chi\kappa} + N_L^{\chi\kappa})^2}{C_L^{\chi\chi} + N_L^{\chi\chi}}. \quad (16)$$

Furthermore, if we decompose  $\kappa$  into uncorrelated pieces sourced by low and high redshifts,  $\kappa = \kappa_{\text{low}} + \kappa_{\text{high}}$ , and assume that  $X(\mathbf{L}) = T(L)\kappa_{\text{low}}(\mathbf{L})$ , so that  $X$  is perfectly correlated with  $\kappa_{\text{low}}$  with a transfer function  $T(L)$ , Eq. (16) becomes

$$C_L^{\hat{\kappa}_{\text{sub}}\hat{\kappa}_{\text{sub}}} = C_L^{\kappa_{\text{low}}\kappa_{\text{low}}} + C_L^{\kappa_{\text{high}}\kappa_{\text{high}}} + N_L^{\kappa\kappa} - \frac{(T(L)C_L^{\kappa_{\text{low}}\kappa_{\text{low}}} + N_L^{\chi\kappa})^2}{T(L)^2 C_L^{\kappa_{\text{low}}\kappa_{\text{low}}} + N_L^{\chi\chi}}. \quad (17)$$

With high noise on  $X$ , we recover  $C_L^{\hat{\kappa}\hat{\kappa}}$ , but in the low-noise limit ( $N_L^{\chi\kappa}, N_L^{\chi\chi} \rightarrow 0$ ), we obtain

$$C_L^{\hat{\kappa}_{\text{sub}}\hat{\kappa}_{\text{sub}}} = C_L^{\kappa_{\text{high}}\kappa_{\text{high}}} + N_L^{\kappa\kappa}, \quad (18)$$

and therefore the low- $z$  contribution to the lensing power spectrum is perfectly subtracted.<sup>6</sup>

In reality, for a tracer we can directly measure, such as cosmic shear, the assumption of perfect and exclusive correlation with the low- $z$  contribution to CMB lensing does not exactly hold; however, if this correlation is sufficiently high, we expect that we can still subtract a significant portion of the unwanted low- $z$  contribution to a lensing map. We consider an explicit example in the following sections.

## B. Cosmic shear from the Rubin Observatory

Cosmic shear is an ideal candidate for an external tracer  $\hat{X}$  that we can use to isolate and subtract the low-redshift contribution to CMB lensing maps. As light rays from distant galaxies travel through the Universe, their paths are deflected by the intervening matter (just as the CMB is lensed), and this introduces correlated ellipticities in the observed images of these galaxies. These correlations, either amongst these galaxies or between the galaxies and another tracer of large-scale structure, are most commonly measured directly from catalogs of observed galaxy shapes. In our forecasts, we will assume that these

<sup>6</sup>An alternative strategy to the map-level subtraction we have considered here would be to perform the forecasts when including all auto- and cross-power spectra between the lensing map and the other tracer, as performed in, e.g., Refs. [64–66] in the context of using CMB lensing to mitigate systematic effects seen in cosmic shear surveys. However, this method is more sensitive to the baryonic biases sourced at low  $z$ , and should be implemented with a mechanism for marginalizing over baryonic models, which we do not consider in this section for simplicity (although we will consider such a mechanism for the CMB lensing alone case in Sec. VI).

catalogs can be converted into lensing convergence maps,<sup>7</sup> and take  $\hat{X}$  to refer to such a map constructed from galaxies in a given redshift bin.

The lensing efficiency for a source galaxy at comoving distance  $\chi_S$  is given by Eq. (2) with  $\chi_{\text{CMB}}$  replaced by  $\chi_S$ :

$$W^\kappa(\chi, \chi_S) = \frac{3}{2} \Omega_m \left( \frac{H_0}{c} \right)^2 \frac{\chi}{a(\chi)} \frac{\chi_S - \chi}{\chi_S}. \quad (19)$$

In practice, galaxies are binned into photometric redshift bins with a finite extent in redshift space, and so we measure the cosmic shear from galaxies at a range of sources. To calculate the shear efficiency for such a bin of galaxies, we integrate over the redshift extent of the bin and weight by the galaxy distribution  $\frac{dn}{d\chi}$ :

$$W^i(\chi) = \frac{1}{n_i} \int_{\chi_i^l}^{\chi_i^f} d\chi_S \frac{dn}{d\chi_S} W^\kappa(\chi, \chi_S), \quad (20)$$

for a bin  $i$  between  $\chi_i^l$  and  $\chi_i^f$ , where  $n_i$  is the total number density of the bin  $n_i = \int_{\chi_i^l}^{\chi_i^f} d\chi \frac{dn}{d\chi}$ . We consider explicitly the distribution predicted for the Rubin Observatory's LSST (Legacy Survey of Space and Time) Gold sample of galaxies [72], a sample that will be used to measure cosmic shear. We take the distribution from Ref. [73]

$$\frac{dn}{dz} = n_s z^{1.24} \exp \left[ - \left( \frac{z}{0.5} \right)^{1.01} \right], \quad (21)$$

with a total number density  $n_s = 26 \text{ arcmin}^{-2}$ .  $\frac{dn}{d\chi}$  can be found from Eq. (21) by computing  $\frac{dn}{d\chi} = \frac{dn}{dz} \frac{dz}{d\chi}$ .

The shear power spectrum of bin  $i$  is given by

$$C_L^{\kappa_i \kappa_i} = \int_0^{\chi_i^f} d\chi \frac{W^i(\chi)^2}{\chi^2} P_m \left( k = \frac{L+1/2}{\chi}, z \right). \quad (22)$$

We consider a survey with  $N$  source bins containing equal numbers of galaxies for the cosmic shear fields, and as we wish to use them to subtract as much of the low- $z$  contribution to the CMB lensing kernel as possible, we combine them in such a way to maximize their correlation with the CMB lensing convergence. As such, we consider a linear combination of shear fields

$$\hat{X} = \sum_i c_i \hat{X}_i, \quad (23)$$

<sup>7</sup>For recent work on such ‘‘mass-mapping’’ techniques for cosmic shear, see Refs. [67–70], several of which are extensions to the method first presented in Ref. [71]. Alternatively, it may be possible to implement our proposal in Sec. V starting directly from shear catalogs, but we leave this to future work.

where  $\hat{X}_i$  is the convergence map of bin  $i$ . The coefficients  $c_i$  are chosen to maximize the correlation coefficient between  $\hat{X}$  and the CMB lensing potential; we compute them following the linear-algebraic methods of Ref. [74] (see their Appendix).<sup>8</sup> The  $c_i$  that maximizes the correlation coefficient of  $\hat{X}$  with the CMB lensing potential  $r_L \equiv \frac{C_L^{\chi\kappa}}{\sqrt{C_L^{\chi\chi} C_L^{\kappa\kappa}}}$  is

$$c_i(L) = \sum_j (C_L^{\kappa_A \kappa_B})_{ij}^{-1} C_L^{j \kappa_{\text{CMB}}}, \quad (24)$$

where  $C_L^{\kappa_A \kappa_B}$  is the covariance matrix of the cosmic shear fields (including noise), and  $C_L^{j \kappa_{\text{CMB}}}$  is the cross-power spectrum between the cosmic shear field  $j$  and the CMB lensing convergence. The elements of  $C_L^{\kappa_A \kappa_B}$  are

$$C_L^{\kappa_A \kappa_B}_{ij} = \int d\chi \frac{W^i(\chi) W^j(\chi)}{\chi^2} P_m \left( k = \frac{L+1/2}{\chi}, z \right) + N_L^{\kappa_i \kappa_j}, \quad (25)$$

with shear noise power spectrum  $N_L^{\kappa_i \kappa_j}$  given by

$$N_L^{\kappa_i \kappa_j} = \delta_{ij} \frac{\sigma_\epsilon^2}{n_i}, \quad (26)$$

where  $\sigma_\epsilon$  is the intrinsic shape noise (we take  $\sigma_\epsilon = 0.26$  [73]) and  $n_i$  is the total angular number density of bin  $i$ . The cross-power spectrum between CMB lensing convergence and cosmic shear in bin  $i$  is given by

$$C_L^{\kappa_j \kappa_{\text{CMB}}} = \int d\chi \frac{W^j(\chi) W_{\text{CMB}}^\kappa(\chi)}{\chi^2} P_m \left( k = \frac{L+1/2}{\chi}, z \right), \quad (27)$$

where  $W_{\text{CMB}}^\kappa(\chi)$  is the CMB lensing efficiency kernel given in Eq. (2).

### C. Intrinsic alignments

It is a nontrivial exercise to separate the apparent ellipticities induced by cosmic shear from the inherent ellipticities of galaxies. Under the assumption that the galaxies have a random distribution of ellipticities, this is not a problem, as taking a high number of galaxies in the sample ensures that the average intrinsic alignment (IA) averages to zero and there is no bias to the signal (although Poissonian noise remains). However, if there is an intrinsic alignment to the galaxies' true ellipticities (e.g., as caused by alignment with the large-scale tidal field), this will bias

<sup>8</sup>Such a map, combining samples in the mid-infrared and far-infrared from *WISE* and *Planck* to be maximally correlated with CMB lensing, was generated in Ref. [75].



one's inference of a lensing signal. With this in mind, we introduce an intrinsic alignment contribution to our forecasts involving cosmic shear to ensure that the effect of intrinsic alignments on the ability of cosmic shear to subtract the low- $z$  information of the CMB lensing kernel is accounted for.

To quantify the intrinsic alignment contribution, the observed ellipticity  $\gamma$  can be separated into a part induced by gravity  $\gamma^G$  and a part that is intrinsic  $\gamma^I$ ,

$$\gamma = \gamma^G + \gamma^I, \quad (28)$$

such that the two-point correlation of  $\gamma$  with itself is

$$\langle \gamma\gamma \rangle = \langle \gamma^G\gamma^G \rangle + \langle \gamma^G\gamma^I \rangle + \langle \gamma^I\gamma^G \rangle + \langle \gamma^I\gamma^I \rangle. \quad (29)$$

In terms of the angular power spectra between the ellipticities of galaxies in two redshift bins labeled by  $i$  and  $j$ , we can write

$$C_L^{\gamma^i\gamma^j} = C_L^{G_iG_j} + C_L^{G_iI_j} + C_L^{G_jI_i} + C_L^{I_iI_j}. \quad (30)$$

Correlations between the tidal field responsible for the intrinsic alignments of foreground galaxies and the gravitational field lensing the images of distant galaxies can cause  $C_L^{G_iI_j}$  to be nonzero when the redshift bin  $j$  is in front of bin  $i$  (negligibly small contributions, which are exactly zero in the Limber approximation, come in the case when  $i$  is in front of  $j$ ). This extra contribution to the correlation is also present in the CMB lensing-cosmic shear cross power spectra:

$$C_L^{\kappa_j\kappa_{\text{CMB}}} = C_L^{G_j\kappa_{\text{CMB}}} + C_L^{I_j\kappa_{\text{CMB}}}. \quad (31)$$

The gravitational contributions  $C_\ell^{G_iG_j}$  and  $C_\ell^{G_j\kappa_{\text{CMB}}}$  are the cosmic shear expressions given in Eqs. (25) and (27). For the other terms, the power spectrum of intrinsic shear  $P_{II}(k, z)$  must be introduced. Then, the  $I - I$  correlations in redshift bin  $i$  are given (within the Limber approximation) by an integral over the redshift bin, weighted by the galaxy density:

$$C_L^{I_iI_i} = \frac{1}{n_i^2} \int_{\chi_j^f}^{\chi_i^f} d\chi \left( \frac{dn}{d\chi_S} \right)^2 P_{II} \left( k = \frac{L+1/2}{\chi}, z \right). \quad (32)$$

The cross-power spectra between bins  $C_L^{I_iI_j}$  is zero (within the Limber approximation) for  $i \neq j$ , as the redshift window functions do not overlap. For the cross-term  $C_L^{GI}$  we introduce the cross-power spectrum between matter and intrinsic ellipticities  $P_{I,m}(k, z)$  such that

$$C_L^{G_iI_j} = \frac{1}{n_j} \int_{\chi_j^f}^{\chi_i^f} d\chi \left( \frac{dn}{d\chi_S} \right) \frac{W^i(\chi)}{\chi^2} P_{I,m} \left( k = \frac{L+1/2}{\chi}, z \right), \quad (33)$$

where  $W^i(\chi)$  can be replaced by  $W_{\text{CMB}}^\kappa(\chi)$  to get  $C_L^{I_j\kappa_{\text{CMB}}}$ . As mentioned earlier, this is only nonzero for  $j$  in front of  $i$  as  $W^i(\chi)$  is zero for  $\chi > \chi_i^f$ , i.e., over the integration range of Eq. (33) if bin  $j$  is behind bin  $i$ .

A simple model for the power spectrum  $P_{II}(k, z)$  assumes that (on large scales) galaxies are aligned with their host dark matter halos, which are given ellipticities by their local tidal field, implying that, in the linear regime [76,77],

$$P_{II}^{\text{lin}}(k, z) = \left( \frac{A_{IA} C_1 \bar{\rho} a^2}{\bar{D}} \right)^2 P_m^{\text{lin}}(k, z), \quad (34)$$

where  $\bar{\rho}$  is the mean matter density of the Universe at redshift  $z$ ,  $\bar{D} = \frac{D(z)}{a}$  where  $D(z)$  is the growth factor normalized to 1 today; and  $P_m^{\text{lin}}(k, z)$  is the linear matter power spectrum.  $C_1$  is a constant amplitude  $C_1 = 5 \times 10^{-14} h^{-2} M_\odot^{-1} \text{Mpc}^3$ , chosen to match the IA amplitude of superCOSMOS in [78] and  $A_{IA}$  is an overall normalization parameter with a fiducial value of unity which should be marginalized over due to uncertainty in the overall amplitude of the intrinsic alignment power. The cross spectrum between the intrinsic ellipticities and the matter power spectrum  $P_{I,m}(k, z)$  is given (again in the linear regime) by

$$P_{I,m}^{\text{lin}}(k, z) = - \left( \frac{A_{IA} C_1 \bar{\rho} a^2}{\bar{D}} \right) P_m^{\text{lin}}(k, z). \quad (35)$$

In [77] this linear alignment model was extended to a ‘‘nonlinear linear alignment model’’ by replacing  $P_m^{\text{lin}}(k, z)$  with the nonlinear  $P_m(k, z)$ . We use this latter model in our forecasts.

In cosmic shear analyses, it is customary to use a very wide prior on the normalization  $A_{IA}$ . However, recent shear surveys have been successful in constraining this parameter in combination with photometric or spectroscopic galaxy clustering: for example, when ignoring galaxy colors, the Dark Energy Survey obtained  $\sigma(A_{IA}) \approx 0.33$  with its year-1 data [79], and the KiDS survey achieved  $\sigma(A_{IA}) \approx 0.5$  in Ref. [80], with both surveys also able to distinguish the amount of intrinsic alignments associated with red and blue galaxies (see also Ref. [81]). Future shear surveys are expected to improve upon this, especially if accompanied by coordinated wide-field spectroscopy (e.g., [82]). Thus, in our baseline forecasts we include a prior of 25% of the fiducial value of  $A_{IA}$ , but we also explore the dependence of our results on this choice.




FIG. 4. The CMB redshift kernel  $W_{\text{CMB}}^{\kappa}$  (black), along with the subtracted field for the  $N_{\text{bins}} = 4$  case (blue). The separate redshift kernels for each bin are also shown in dashed lines. The coefficient of each bin in the linear combination is included in the caption.

#### D. Impact of shear subtraction on neutrino mass inference

Subtracting the optimally-combined shear map of Sec. VB from the CMB lensing map results in a significant reduction of the baryonic bias on  $M_{\nu}$  without an appreciable increase in the statistical uncertainty. To illustrate the redshift distribution of the field we subtract, we show in

Fig. 4 the redshift kernel  $W(\chi) \frac{dz}{d\chi}$  for CMB lensing, cosmic shear for the case when  $N_{\text{bins}} = 4$ , and the linear combination of the 4 shear bins as determined by Eq. (23). We also quote the numerical values of the relevant  $c_i$ , which demonstrate that the highest-redshift shear bin is the dominant contribution to the optimal combination, but with the other bins also playing a role in improving the correlation with the CMB lensing redshift kernel.

In Fig. 5, we plot the constraints and bias on  $M_{\nu}$  against  $N_{\text{bins}}$ , the number of galaxy redshift bins we use to construct  $\hat{X}$ . Increasing  $N_{\text{bins}}$  does not increase the number of galaxies in the analysis, but it does have the effect of increasing our redshift resolution and allowing us to weight the galaxy kernels to match the CMB lensing kernel better. Note, however, that at some value of  $N_{\text{bins}}$ , photometric redshift errors will not allow for increasing  $N_{\text{bins}}$  to give us better redshift resolution, and we expect the curves to saturate around this point in a treatment where photometric redshifts are included. However, for the redshift errors expected of the Rubin Observatory's LSST, we expect to be able to significantly subtract the bias on  $M_{\nu}$  with this method even with a small number of bins, in which case photometric redshift errors will likely be insignificant due to the wide extents of the bins.

Figure 5 shows that, with  $L_{\text{max}} = 3100$  for CMB lensing, implementing the shear subtraction generally decreases the




FIG. 5. Demonstrating mitigation strategy 2: the decrease in the bias when an optimal combination of  $N_{\text{bins}}$  cosmic shear bins is subtracted from the CMB lensing map. With the exception of  $N_{\text{bins}} = 0$  (which corresponds to the no-subtraction case), the total redshift extent and number of galaxies are held constant, with  $N_{\text{bins}}$  controlling the slice thickness. On top is for an analysis with  $L_{\text{max}}$  of 3100; on bottom we have applied a scale cut of  $L_{\text{max}} = 1000$  (as well as subtracting cosmic shear). All curves shown include marginalization over intrinsic alignments for cosmic shear, with a 25% prior on the intrinsic alignment amplitude. It is clear that the bias on  $M_{\nu}$  can be reduced significantly in the latter case, without appreciably loosening the constraint on  $M_{\nu}$ .

bias on  $M_\nu$  associated with all but one simulation by factors of 3 to 7 for SO, and a factor of  $\sim 2$  for S4 (we denote the no-subtraction case by  $N_{\text{bins}} = 0$  in the figure). With  $L_{\text{max}} = 1000$ , we similar improvements for SO, but closer to a factor of 5 improvement for S4, in addition to the factor of 2 improvement arising from the  $L_{\text{max}}$  cut. In all cases, the bulk of the improvement can be obtained with only 2 shear redshift bins (with boundaries  $z = [0, 0.93, 4]$ ) for the LSST specifications we use. Concurrent with the decrease in bias is an increase in the statistical uncertainty on  $M_\nu$ , by at most  $\sim 20\%$ .

Most simulations in Fig. 5 display similar behavior, with the bias on  $M_\nu$  decreasing with  $N_{\text{bins}}$ , but Horizon-AGN exhibits the opposite trend, with the bias *increasing* as more redshift bins are used for the subtraction. This is due to the stronger effect of baryons on high-redshift ( $z \gtrsim 3$ ) clustering observed in Horizon-AGN as compared to the other simulations (contrast Fig. 2 of Ref. [48] with results from other simulations summarized in Refs. [43,55]). If the high-redshift clustering in Horizon-AGN (primarily driven by gas pressure delaying the collapse of dark matter into halos, rather than AGN feedback [48]), as opposed to that in the other simulations (in which the effect of baryons at high  $z$  is much more mild) is reproduced in the actual Universe, then isolating the high- $z$  part of a CMB lensing map will not be sufficient to mitigate the impact of baryons on a neutrino mass constraint.

In Fig. 6, we show the effect of varying the prior on the intrinsic alignment amplitude  $A_{IA}$  on the uncertainty and bias on  $M_\nu$ , for  $N_{\text{bins}} = 5$  and the bias associated with the OWLS-AGN simulation. For CMB-S4 with the tighter  $\tau$  prior, loosening  $\sigma(A_{IA})$  from 0.25 to 3 degrades the expected error bar on  $M_\nu$  by 15% for  $L_{\text{max}} = 3100$  or 40% for  $L_{\text{max}} = 1000$ , and increases the bias on  $M_\nu$  by a factor of  $\sim 2$  compared to the  $\sigma(A_{IA}) = 0.25$  case. Thus, with the wider  $A_{IA}$  prior, the shear subtraction combined

with the  $L_{\text{max}}$  cut is generally still able to reduce the bias on  $M_\nu$  by a factor of  $\sim 5$ –6 compared to the case with no mitigation strategy. On the other hand, exact knowledge of  $A_{IA}$  leads to improvements in the uncertainty and bias on  $M_\nu$  of a few tens of percents. Note that the nonlinear alignment model we have used for intrinsic alignments will likely be superseded by more detailed models for future shear surveys (e.g., [83,84]), but our exploration of priors on  $A_{IA}$  can be taken as indicative of how one’s knowledge of intrinsic alignments affects the shear subtraction procedure we have described.

## VI. STRATEGY 3: MARGINALIZATION OVER ADDITIONAL PARAMETERS

The mitigation strategies we considered above involved removing portions of the data most sensitive to baryons, via scale cuts and/or subtracting proxies for low-redshift information in CMB lensing, while retaining as much constraining power on  $M_\nu$  as possible. However, as biased constraints on  $M_\nu$  arise from neglecting baryonic effects in the theoretical modelling of the matter power spectrum, one can instead incorporate a model for these effects; by marginalizing over the associated parameters, we can hope to reduce the bias on  $M_\nu$  without requiring precise knowledge of the impact of baryons. Examples of such models include perturbation theory [58,85,86], extended halo models [32,56,87–90], empirical fitting functions [20,43], or principal-component decompositions [91–93] for the matter power spectrum from simulations, “baryonification” algorithms that modify the outputs of  $N$ -body simulations [94–97], emulators [98], or approaches based on machine learning [99,100].

We choose to test this marginalization approach using the model from Ref. [32]. This model is based on the halo model (e.g., [101]), with a modified 1-halo term for the




FIG. 6. The effect of marginalizing over the intrinsic alignment amplitude. If a prior is not included the constraint on  $\sigma(M_\nu)$  can increase significantly, as is shown on the left; the bias also increases, as is shown on the right for OWLS-AGN.

TABLE I. Expected uncertainty and bias on  $M_\nu$  with and without marginalization over the  $A$  and  $\eta_0$  parameters of modified halo model from Ref. [32] (denoted ‘‘After marg.’’ and ‘‘No marg.’’ respectively). Expected uncertainty and bias on  $M_\nu$  with and without marginalization over the  $A$  and  $\eta_0$  parameters of modified halo model from Ref. [32], using CMB lensing multipoles up to  $L_{\max} = 3100$ . With marginalization, the biases are reduced by a factor of  $\sim 10$  in most cases, while the uncertainty is only degraded by  $\sim 15\%$  at most.

Experiment	$\sigma_{\text{prior}}(\tau)$	$\sigma(M_\nu)$ [meV]		Simulation	Bias $\Delta M_\nu$ [meV]		$\Delta M_\nu / \sigma(M_\nu)$	
		No marg.	After marg.		No marg.	After marg.	No marg.	After marg.
SO	0.006	22	24	OWLS-AGN	7.2	0.14	0.32	0.0060
				BAHAMAS	6.6	0.74	0.29	0.031
				BAHAMAS-LowAGN	3.8	0.43	0.17	0.018
				BAHAMAS-HighAGN	12	1.4	0.55	0.059
				Horizon-AGN	0.88	-0.81	0.039	-0.033
				TNG100	0.92	0.090	0.041	0.0038
				TNG300	0.92	0.088	0.042	0.0037
S4	0.006	20	22	OWLS-AGN	16	0.39	0.76	0.018
				BAHAMAS	13	1.1	0.65	0.052
				BAHAMAS-LowAGN	7.7	0.67	0.38	0.031
				BAHAMAS-HighAGN	24	2.4	1.2	0.11
				Horizon-AGN	3.2	-0.78	0.16	-0.036
				TNG100	2.0	0.22	0.098	0.010
				TNG300	2.1	0.20	0.10	0.0091
S4	0.002	12	14	OWLS-AGN	14	0.37	1.2	0.027
				BAHAMAS	12	1.2	1.0	0.083
				BAHAMAS-LowAGN	7.2	0.68	0.59	0.049
				BAHAMAS-HighAGN	22	2.4	1.8	0.18
				Horizon-AGN	2.8	-0.89	0.23	-0.064
				TNG100	1.9	0.21	0.15	0.015
				TNG300	2.0	0.19	0.16	0.014

matter power spectrum. Baryonic effects are parametrized by two parameters. The first,  $A$ , is the amplitude of the halo concentration-mass relationship  $c(M, z)$ :

$$c(M, z) = A \frac{1 + z_f(M)}{1 + z}, \quad (36)$$

where  $z_f(M)$  is the formation redshift of halos of mass  $M$ ; this is designed to capture the effects of processes such as gas cooling, which can cause increased halo concentration. The second parameter,  $\eta$ , alters the (Fourier-transformed) halo density profile  $u(k, M)$  via

$$u(k, M) \rightarrow u(\nu^\eta k, M), \quad (37)$$

where  $\nu = \frac{\delta_c}{\sigma(M)}$  with  $\delta_c$  the critical density required for spherical collapse and  $\sigma(M)$  variance in the initial density fluctuation field when smoothed with a top-hat filter with the size of the virial radius of the halo. The Fourier-transformed halo profile is given by

$$u(k, M) \equiv \frac{1}{M} \int_0^{r_v} dr 4\pi r^2 \rho(r, M) \frac{\sin(kr)}{kr}, \quad (38)$$

where  $r_v$  is the virial radius of the halo and  $\rho(r, M)$  is the halo density profile in real space. For positive  $\eta$ , the modification (37) ‘‘puffs out’’ higher-mass halos ( $\nu > 1$ )

and contracts lower mass halos, and as such  $\eta$  is referred to as the ‘‘halo bloating parameter.’’ This is intended to capture some of the effects of AGN feedback on halo profiles. By fitting  $A$  and  $\eta$  to the OWLS simulations, Ref. [56] found that it was adequate to use a single redshift-independent value for  $A$ , while the redshift dependence of  $\eta$  was well captured by

$$\eta(z) = \eta_0 - 0.3\sigma_8(z), \quad (39)$$

with a single  $\eta_0$ , where  $\sigma_8$  is the variance of density fluctuations over a sphere with radius  $8h \text{ Mpc}^{-1}$ , and we use those choices in our calculations.

To ensure that the parameter space of  $A$  and  $\eta_0$  sufficiently describes the baryonic effects in the simulations we are considering, in Appendix we fit these parameters to the  $P_{\text{bary}}/P_{\text{DMO}}$  ratios from each simulation, and compare the ‘‘best-fit’’ predictions with the simulations’ measurements. We indeed find that the model is able to reproduce all simulation results with a precision of  $\sim 5\%$  over the scales we are concerned with, which is an acceptable level since we are more concerned with the range of simulation results rather than exactly reproducing any one simulation. This agreement also justifies our use of the model from Ref. [32] as opposed to turning to more recent updates (e.g., [88,89]).






FIG. 7. The derivatives of the lensing power spectrum with respect to the parameters of the model for baryonic effects from Ref. [32], and also  $M_\nu$ , all normalized by their own values at  $L = 1000$ .  $M_\nu$  is tending towards a scale-independent effect on  $C_L^{KK}$  across a wide range of scales, while  $A$  and  $\eta_0$  are significantly scale dependent, becoming increasingly important on small scales.

To include this model in our forecasts, we use the same formalism as in Sec. III, but expanding the vector of parameters in Eq. (3) to include  $A$  and  $\eta_0$ :

$$\vec{\theta} = (h, \Omega_b h^2, \Omega_c h^2, \tau, n_s, A_s, M_\nu; A, \eta_0). \quad (40)$$

The fiducial values we use for  $A$  and  $\eta_0$  are  $A = 3.13$  and  $\eta_0 = 0.60$ , corresponding to those fit to the DMO run of OWLS in Ref. [56]. We do not assume any prior knowledge of the true values of these parameters, which is a conservative choice given the multitude of other datasets which could likely constrain them at some level. Upon marginalizing over  $A$  and  $\eta_0$ , along with the other cosmological parameters, we find the results in Table I. In particular, we find that the uncertainty on  $M_\nu$  is degraded by only  $\sim 10\%$  compared to the case where baryonic effects are ignored in the modeling (equivalent to fixing  $A$  and  $\eta_0$  to their fiducial values), while the biases on  $M_\nu$  are drastically reduced, by factors of  $\sim 10$  or more. In particular, for CMB-S4 with the tightest  $\tau$  prior, the bias corresponding to BAHAMAS-HighAGN is roughly  $0.2\sigma$ , while for all other simulations it is less than  $0.1\sigma$ ; without marginalization, there are three simulations that induce a bias exceeding  $1\sigma$ .

To understand why this prescription works so well at removing the bias while preserving the constraining power, it is helpful to plot the derivatives of  $C_L^{KK}$  with respect to  $A$ ,  $\eta_0$ , and  $M_\nu$ ; see Fig. 7. We see that  $A$  and  $\eta_0$  have significantly different effects on the shape of the lensing power spectrum than  $M_\nu$ : since they only modify the 1-halo term in the matter power spectrum, they have the strongest impact at small scales, while neutrino mass suppresses

structure growth over a wider range of scales (recall Fig. 1). This lack of degeneracy implies that the  $M_\nu$  constraint is not degraded when  $A$  and  $\eta_0$  are marginalized over; furthermore, since the model covers the space of baryonic effects well, the marginalization is effective at removing the associated bias from a determination of  $M_\nu$ . These conclusions are consistent with other studies of cosmic shear [17,18,102], which have found that when priors from the primary CMB or other observations are included, marginalizing over a baryonic model with only a few parameters enables unbiased constraints on neutrino mass without large increases in uncertainty.

## VII. DISCUSSION AND CONCLUSION

Upcoming measurements of CMB lensing have great promise to measure the sum of neutrino masses ( $M_\nu$ ), but this measurement will only be possible if each of several systematic effects are tightly controlled. In this work, we considered one such systematic, related to the impact of ‘‘baryonic effects’’ (the name given to astrophysical processes like gas cooling and AGN feedback) on the lensing power spectrum. Recent simulations indicate that uncertainty in these effects can bias a neutrino mass measurement from CMB lensing by a sizeable fraction of the statistical error bar if they are not incorporated in the modeling or mitigated in some other way [19].

We investigated three strategies for mitigating this bias, using Fisher forecasts that combine expected CMB lensing and primary CMB measurements from upcoming experiments, BAO constraints expected from DESI, and a prior on the mean optical depth  $\tau$ , either from current Planck constraints or assuming a cosmic-variance-limited measurement. Our chosen strategies and results are as follows:

- (1) Decreasing the largest multipole  $L_{\max}$  (smallest physical distance) used for neutrino mass constraints (Sec. IV): baryonic effects are stronger at smaller scales, so one can attempt to reduce sensitivity to these effects by excluding smaller scales from an analysis. We found that keeping  $L_{\max} \gtrsim 1000$  is necessary to avoid sacrificing significant constraining power on  $M_\nu$ , and that setting  $L_{\max} = 1000$  can reduce the bias on  $M_\nu$  by as much as a factor of 2 (compared to our fiducial forecast with  $L_{\max} = 3100$ ). However, this reduction still allows the bias be of the same order as the statistical error bar on  $M_\nu$ , so further mitigation is needed.
- (2) Removing the low-redshift contribution to the lensing map using external tracers (Sec. V): CMB lensing probes structures over a wide redshift range, while simulations indicate that baryonic effects should only have a sizeable effect on clustering at the lower end of this range. Thus, one can attempt to remove the low- $z$  contribution from a CMB lensing map, and use the resulting map for a neutrino mass analysis. We have implemented this proposal using a




FIG. 8. A summary of the efficacy of our mitigation methods. For each experimental configuration and each mitigation method, we show the forecast error bars for  $M_\nu$ , centered on the fiducial value of 60 meV but offset by the bias forecast for each simulation. Without any mitigation, the biases are substantial and have differing values for different simulations [19]. However, it can be seen that  $L_{\max}$  reduction and shear subtraction each reduce the biases to some extent, and that when both methods are performed the biases are almost completely removed. Marginalization over baryonic parameters also does an excellent job of almost completely removing the biases.

weighted sum of cosmic shear measurements in different source redshift bins. We found that, with  $L_{\max} = 1000$ , subtracting a combination of 2 shear redshift bins (modeled on the Rubin Observatory’s LSST) from a CMB lensing map is sufficient to reduce the bias on  $M_\nu$  by factors of 5 or more, down to 3 meV for the most extreme simulation we consider. We also considered the effect of intrinsic alignments in the LSST-like data on this procedure, and found a minor impact, when assuming the priors that are expected.

- (3) Marginalizing over a parametrization of baryonic effects (Sec. VI): if the effects of neutrino mass and baryonic processes are sufficiently nondegenerate in the lensing power spectrum, then one can consider including a simple parametrization of the latter in the matter power spectrum model, and marginalizing over the corresponding parameters. After checking that the modified halo model from Ref. [32] can describe our set of simulations with appropriate accuracy, we found that marginalizing over the model’s two parameters only degrades the constraining power on  $M_\nu$  by about 15%, while reducing

the bias on  $M_\nu$  by factors of 10 or more, down to  $\sim 2$  meV for the most extreme simulation. This is true even without a reduction in  $L_{\max}$ .

We conclude that either of strategies 2 or 3 should be sufficient to reduce the bias in  $M_\nu$  from baryons to an acceptable level ( $\Delta M_\nu \lesssim 0.2\sigma(M_\nu)$ ) for CMB-S4. These results are summarized in Fig. 8.

Methods 2 and 3 are complementary in that method 2 relies on combining external data with CMB lensing maps, while method 3 relies on the inclusion of extra modeling ingredients. They can serve as useful cross-checks of each other: the same result from two different approaches would indicate a robust measurement of  $M_\nu$ .

There is a fourth mitigation strategy that we have not explored in this work: using external measurements to constrain or fix the form of baryonic effects on CMB lensing. It has been shown that the matter power spectrum suppression seen in a range of hydrodynamical simulations is strongly correlated with the mean baryon fraction of group- and cluster-scale halos ( $M_{\text{halo}} \sim 10^{14} h^{-1} M_\odot$ ) [43,90], so an external constraint on this quantity would likely give a much sharper picture of how much power suppression to include in the modeling. This constraint may

be achievable with future x-ray observations of groups and clusters (e.g., [103]). For example, when using the model for baryonic effects from Ref. [95], Ref. [18] finds that gas fraction measurements by the upcoming eROSITA telescope [104], combined with cluster mass estimates from Euclid weak lensing, will significantly reduce the uncertainty on  $M_\nu$  from a combination of cosmic shear and Planck CMB measurements. However, inference of gas fractions from x-ray measurements requires a measure of the total halo mass, and the various methods to obtain this (e.g., assuming hydrostatic equilibrium or using weak lensing of background galaxies) each come with their own caveats. Furthermore, x-ray measurements are most sensitive to hot gas in a halo's interior, while Refs. [90,94] have shown that gas at the outskirts of groups and clusters has an important effect on the matter power spectrum.

This fainter, more diffuse gas can be probed using the thermal and kinetic Sunyaev-Zel'dovich effects. Mean gas fractions can be extracted from cross-correlations between CMB maps and group/cluster catalogs [105–108], or more generally, stacked gas and pressure profiles can be measured directly from these cross-correlations [109–112], with upcoming surveys promising to provide much more powerful measurements [113,114]. Further constraints are possible by correlating thermal Sunyaev-Zel'dovich maps and cosmic shear [115,116] or CMB lensing itself [117], helping to break degeneracies between baryonic effects and neutrino mass. Clearly, there are many avenues for independent constraints of baryonic effects, which can be incorporated into an analysis of CMB lensing. However, the strategies presented in this paper do not depend on any such constraints, and therefore represent a promising approach to pursue in parallel.

Finally, it is worth noting that we have made specific choices when implementing these strategies in our forecasts, but other choices are possible. For instance, one could choose a low-redshift tracer other than cosmic shear to implement the map-level subtraction from Sec. V; another option would be to use spectroscopic or photometric galaxy catalogs, although galaxy bias and selection effects would need to be carefully accounted for in the subtraction procedure. One could also consider marginalizing over other models for baryonic effects, such as effective-field-theory-based perturbation theory, which Ref. [58] found to be capable of describing baryonic effects on CMB lensing at  $L \lesssim 2000$  with suitable accuracy for CMB-S4. Regardless of these specific choices, we expect our general conclusions about these strategies to hold. Therefore, it appears that baryonic effects on a neutrino mass constraint from CMB lensing can be straightforwardly reduced to a negligible level.

#### ACKNOWLEDGMENTS

We thank Mat Madhavacheril and Blake Sherwin for interesting discussions that provided motivation for this

work, and we also thank Shahab Joudaki, Matthew Lewandowski, and Emmanuel Schaan for useful discussions. We thank Francisco Villaescusa-Navarro for measuring and providing IllustrisTNG power spectra used in this work, the OWLS and Horizon-AGN teams for making their power spectra publicly available, Mat Madhavacheril for making his forecasting code publicly available, and Eegene Chung for sharing her lensing and forecasting codes. Research at Perimeter Institute is supported in part by the Government of Canada through the Department of Innovation, Science and Industry Canada and by the Province of Ontario through the Ministry of Colleges and Universities. FMcC acknowledges support from the Vanier Canada Graduate Scholarships program.

#### APPENDIX: COMPARISON OF SIMULATIONS AND PARAMETRIC MODEL FOR BARYONIC EFFECTS

In Sec. VI, we investigated whether including a model for baryonic effects in the prediction for the matter power spectrum can reduce the bias on  $M_\nu$  after marginalization over the model's parameters. For this to be a valid method, we must be reasonably confident that the model can capture a realistic range of baryonic effects. In this appendix, we check this for the hydrodynamical simulations we use in this work, assuming that this set of simulations itself spans a realistic range of effects. More detailed comparisons between these simulations and observations will likely be required to fully justify this assumption, but this is beyond the scope of this work.

We compare model predictions from Ref. [32] with the power spectrum ratios  $\hat{R}(k, z)$  measured from simulations, using the following statistic:

$$\Delta(A, \eta_0) \equiv \sum_{i,j} \left( \hat{R}(k_i, z_j) - \frac{P_m(k_i, z_j; A, \eta_0)}{P_{\text{DMO}}(k_i, z_j)} \right)^2, \quad (\text{A1})$$

where  $P_{\text{DMO}}$  is evaluated with the fiducial values of  $A$  and  $\eta_0$  from Sec. VI. We sum over  $k$  and  $z$  points at which simulation measurements are available over  $0 \leq z \leq 2$  and  $1 \text{ Mpc}^{-1} \leq k \leq 10 \text{ Mpc}^{-1}$ , chosen to correspond roughly to the ranges in which baryonic effects significantly affect  $C_L^{\kappa\kappa}$  for  $L \lesssim 3000$  [19]. Equation (A1) is equivalent to a  $\chi^2$  statistic that weights all points equally, motivated by other work that has found sample-variance uncertainties on  $\hat{R}$  to be roughly scale independent [43,48,55]. Our goal is only to examine the best-fit predictions of the model for each simulation, rather than fully quantify the goodness of fit (which we cannot do without better knowledge of the uncertainties on  $\hat{R}$ ), so we simply use unit weights in Eq. (A1). We evaluate  $P_m$  and  $P_{\text{DMO}}$  at our fiducial cosmology, because our goal is to check how well the model from Ref. [32] can reproduce our range of  $\hat{R}$  curves with cosmology held fixed.




FIG. 9. Contour plots in the  $A, \eta_0$  plane. We show contours of  $\Delta(A, \eta_0)$ , where  $\Delta(A, \eta_0)$  is defined in Eq. (A1). The points which minimize  $\Delta(A, \eta_0)$  are shown, and the contours are filled in for values of  $\Delta(A, \eta_0)$  which  $\Delta(A, \eta_0) < \Delta_{\min} + 5 \times 10^{-5}$ , corresponding to the separate  $\Delta_{\min}$  for each simulation.

In Fig. 9, we show contour plots of  $\Delta(A, \eta_0)$  for each simulation. These plots clearly imply a degeneracy between  $A$  and  $\eta_0$  for each simulation, roughly consistent with the degeneracy directions seen in the fits in Ref. [56] (see their Fig. 6). Figure 10 shows the best-fitting

predictions for  $\hat{R}(k, z)$  from minimizing  $\Delta(A, \eta_0)$  with respect to the two parameters, at a few representative redshifts. We find that for all simulations, the model can describe the power spectrum suppression to better than  $\sim 5\%$  over the scales of interest, with better fits at




FIG. 10. Top: the response functions  $\hat{R}(k, z)$  measured from the simulations (solid lines) and the  $R(A, \eta_0; k, z)$  (from the baryonic model) that minimize  $\Delta(A, \eta_0)$ , at various redshifts. Bottom: ratios of best-fit and measured  $R$  functions. A dashed vertical line is shown at the  $k$  that is equivalent to  $3100/\chi(z)$  at each redshift, i.e., the maximum  $k$  used to calculate  $C_L^{\kappa}$  in the Limber approximation at each  $z$ .






FIG. 11. Left: the ratio of the baryonic lensing power spectra to the DMO power spectra, computed with the  $P(k, z; A, \eta_0)$  at the “best fit” values found for each simulation (in dashed lines). Also shown (in solid lines) is the ratio of the power spectra with the response function  $\hat{R}(k, z)$  measured directly from the simulations. Right: the ratio of the lensing power spectra computed at the best-fit  $A, \eta_0$  and that computed with  $\hat{R}(k, z)$  from the simulations. We find that the model is capable of reproducing the simulation measurements at better than 2% accuracy for  $L < 3100$ .

lower  $z$ . While other models have been shown to match a subset of these simulations at higher precision (e.g., [89]), the  $\sim 5\%$  precision we find for the model from Ref. [32] is sufficient to use it our proof-of-concept forecast in Sec. VI. To see this, note that 5% systematic errors in

$P_m$  over the scales we fit for translate into  $\sim 2\%$  errors in  $C_L^{\kappa\kappa}$  (see Fig. 11), and Fig. 1 shows that the simulation-derived  $C_L^{\kappa\kappa}$  curves are still distinguishable from the effect of massive neutrinos even with this level of errors.

- 
- [1] Y. Fukuda *et al.* (Super-Kamiokande Collaboration), *Phys. Rev. Lett.* **81**, 1562 (1998).
  - [2] Q. R. Ahmad *et al.* (SNO Collaboration), *Phys. Rev. Lett.* **89**, 011301 (2002).
  - [3] C. Patrignani *et al.* (Particle Data Group), *Chin. Phys. C* **40**, 100001 (2016).
  - [4] N. Aghanim *et al.* (Planck Collaboration), *Astron. Astrophys.* **641**, A6 (2020).
  - [5] J. Lesgourgues and S. Pastor, *Phys. Rep.* **429**, 307 (2006).
  - [6] J. Lesgourgues and S. Pastor, *Adv. High Energy Phys.* **2012**, 608515 (2012).
  - [7] P. Ade *et al.* (Simons Observatory Collaboration), *J. Cosmol. Astropart. Phys.* **02** (2019) 056.
  - [8] B. Benson *et al.* (SPT-3G Collaboration), *Proc. SPIE Int. Soc. Opt. Eng.* **9153**, 91531P (2014).
  - [9] K. N. Abazajian *et al.* (CMB-S4 Collaboration), arXiv: 1610.02743.
  - [10] R. Allison, P. Caucal, E. Calabrese, J. Dunkley, and T. Louis, *Phys. Rev. D* **92**, 123535 (2015).
  - [11] M. J. White, *Astropart. Phys.* **22**, 211 (2004).
  - [12] H. Zhan and L. Knox, *Astrophys. J.* **616**, L75 (2004).
  - [13] Y. P. Jing, P. Zhang, W. P. Lin, L. Gao, and V. Springel, *Astrophys. J.* **640**, L119 (2006).
  - [14] D. H. Rudd, A. R. Zentner, and A. V. Kravtsov, *Astrophys. J.* **672**, 19 (2008).
  - [15] E. Semboloni, H. Hoekstra, J. Schaye, M. P. van Daalen, and I. J. McCarthy, *Mon. Not. R. Astron. Soc.* **417**, 2020 (2011).
  - [16] A. Natarajan, A. R. Zentner, N. Battaglia, and H. Trac, *Phys. Rev. D* **90**, 063516 (2014).
  - [17] D. Copeland, A. Taylor, and A. Hall, *Mon. Not. R. Astron. Soc.* **493**, 1640 (2020).
  - [18] A. Schneider, A. Refregier, S. Grandis, D. Eckert, N. Stoira, T. Kacprzak, M. Knabenhans, J. Stadel, and R. Teyssier, *J. Cosmol. Astropart. Phys.* **04** (2020) 020.
  - [19] E. Chung, S. Foreman, and A. van Engelen, *Phys. Rev. D* **101**, 063534 (2020).
  - [20] J. Harnois-Déraps, L. van Waerbeke, M. Viola, and C. Heymans, *Mon. Not. R. Astron. Soc.* **450**, 1212 (2015).

- [21] S. Foreman, M. R. Becker, and R. H. Wechsler, *Mon. Not. R. Astron. Soc.* **463**, 3326 (2016).
- [22] H.-J. Huang *et al.* (DES Collaboration), *Mon. Not. R. Astron. Soc.* **502**, 6010 (2021).
- [23] M. Yoon and M. J. Jee, *Astrophys. J.* **908**, 13 (2021).
- [24] A. Lewis and A. Challinor, *Phys. Rep.* **429**, 1 (2006).
- [25] D. N. Limber, *Astrophys. J.* **117**, 134 (1953).
- [26] M. LoVerde and N. Afshordi, *Phys. Rev. D* **78**, 123506 (2008).
- [27] G. Pratten and A. Lewis, *J. Cosmol. Astropart. Phys.* **08** (2016) 047.
- [28] G. Fabbian, M. Calabrese, and C. Carbone, *J. Cosmol. Astropart. Phys.* **02** (2018) 050.
- [29] A. R. Duffy, J. Schaye, S. T. Kay, C. Dalla Vecchia, R. A. Battye, and C. Booth, *Mon. Not. R. Astron. Soc.* **405**, 2161 (2010).
- [30] [http://github.com/sjforeman/cmlensing\\_baryons](http://github.com/sjforeman/cmlensing_baryons).
- [31] A. Lewis, A. Challinor, and A. Lasenby, *Astrophys. J.* **538**, 473 (2000).
- [32] A. Mead, C. Heymans, L. Lombriser, J. Peacock, O. Steele, and H. Winther, *Mon. Not. R. Astron. Soc.* **459**, 1468 (2016).
- [33] A. Aghamousa *et al.* (DESI Collaboration), [arXiv:1611.00036](https://arxiv.org/abs/1611.00036).
- [34] J. Peloton, M. Schmittfull, A. Lewis, J. Carron, and O. Zahn, *Phys. Rev. D* **95**, 043508 (2017).
- [35] D. Green, J. Meyers, and A. van Engelen, *J. Cosmol. Astropart. Phys.* **12** (2017) 005.
- [36] A. Cooray, *Phys. Rev. D* **65**, 103510 (2002).
- [37] A. Lewis, A. Challinor, and D. Hanson, *J. Cosmol. Astropart. Phys.* **03** (2011) 018.
- [38] L. Pagano, J. M. Delouis, S. Mottet, J. L. Puget, and L. Vibert, *Astron. Astrophys.* **635**, A99 (2020).
- [39] Z. Li, V. Gluscevic, K. K. Boddy, and M. S. Madhavacheril, *Phys. Rev. D* **98**, 123524 (2018).
- [40] <https://github.com/msyriac/pyfisher>.
- [41] M. P. van Daalen, J. Schaye, C. M. Booth, and C. D. Vecchia, *Mon. Not. R. Astron. Soc.* **415**, 3649 (2011).
- [42] M. P. van Daalen, J. Schaye, C. Booth, and C. D. Vecchia, *Mon. Not. R. Astron. Soc.* **415**, 3649 (2011).
- [43] M. P. van Daalen, I. G. McCarthy, and J. Schaye, *Mon. Not. R. Astron. Soc.* **491**, 2424 (2020).
- [44] I. G. McCarthy, J. Schaye, S. Bird, and A. M. Le Brun, *Mon. Not. R. Astron. Soc.* **465**, 2936 (2017).
- [45] I. G. McCarthy, S. Bird, J. Schaye, J. Harnois-Deraps, A. S. Font, and L. Van Waerbeke, *Mon. Not. R. Astron. Soc.* **476**, 2999 (2018).
- [46] Y. Dubois *et al.*, *Mon. Not. R. Astron. Soc.* **444**, 1453 (2014).
- [47] Y. Dubois *et al.*, *Mon. Not. R. Astron. Soc.* **463**, 3948 (2016).
- [48] N. E. Chisari, M. L. Richardson, J. Devriendt, Y. Dubois, A. Schneider, M. Brun, Amandine Le, R. S. Beckmann, S. Peirani, A. Slyz, and C. Pichon, *Mon. Not. R. Astron. Soc.* **480**, 3962 (2018).
- [49] A. Pillepich *et al.*, *Mon. Not. R. Astron. Soc.* **475**, 648 (2018).
- [50] V. Springel *et al.*, *Mon. Not. R. Astron. Soc.* **475**, 676 (2018).
- [51] D. Nelson *et al.*, *Mon. Not. R. Astron. Soc.* **475**, 624 (2018).
- [52] J. P. Naiman *et al.*, *Mon. Not. R. Astron. Soc.* **477**, 1206 (2018).
- [53] F. Marinacci *et al.*, *Mon. Not. R. Astron. Soc.* **480**, 5113 (2018).
- [54] D. Nelson *et al.*, [arXiv:1812.05609](https://arxiv.org/abs/1812.05609).
- [55] S. Foreman, W. Coulton, F. Villaescusa-Navarro, and A. Barreira, *Mon. Not. R. Astron. Soc.* **498**, 2887 (2020).
- [56] A. Mead, J. Peacock, C. Heymans, S. Joudaki, and A. Heavens, *Mon. Not. R. Astron. Soc.* **454**, 1958 (2015).
- [57] B. O. Mummery, I. G. McCarthy, S. Bird, and J. Schaye, *Mon. Not. R. Astron. Soc.* **471**, 227 (2017).
- [58] D. P. Bragança, M. Lewandowski, D. Sekera, L. Senatore, and R. Sgier, [arXiv:2010.02929](https://arxiv.org/abs/2010.02929).
- [59] K. Abazajian *et al.*, [arXiv:1907.04473](https://arxiv.org/abs/1907.04473).
- [60] T. Okamoto and W. Hu, *Phys. Rev. D* **67**, 083002 (2003).
- [61] A. van Engelen, S. Bhattacharya, N. Sehgal, G. Holder, O. Zahn, and D. Nagai, *Astrophys. J.* **786**, 13 (2014).
- [62] S. J. Osborne, D. Hanson, and O. Doré, *J. Cosmol. Astropart. Phys.* **03** (2014) 024.
- [63] K. M. Smith, D. Hanson, M. LoVerde, C. M. Hirata, and O. Zahn, *J. Cosmol. Astropart. Phys.* **06** (2012) 014.
- [64] S. Das, J. Errard, and D. Spergel, [arXiv:1311.2338](https://arxiv.org/abs/1311.2338).
- [65] E. Schaan, E. Krause, T. Eifler, O. Doré, H. Miyatake, J. Rhodes, and D. N. Spergel, *Phys. Rev. D* **95**, 123512 (2017).
- [66] E. Schaan, S. Ferraro, and U. Seljak, *J. Cosmol. Astropart. Phys.* **12** (2020) 001.
- [67] N. Jeffrey, F. Lanusse, O. Lahav, and J.-L. Starck, *Mon. Not. R. Astron. Soc.* **492**, 5023 (2020).
- [68] B. Mawdsley *et al.* (DES Collaboration), *Mon. Not. R. Astron. Soc.* **493**, 5662 (2020).
- [69] S. Pires *et al.* (EUCLID Collaboration), *Astron. Astrophys.* **638**, A141 (2020).
- [70] M. A. Price, J. D. McEwen, L. Pratley, and T. D. Kitching, *Mon. Not. R. Astron. Soc.* **500**, 5436 (2020).
- [71] N. Kaiser and G. Squires, *Astrophys. J.* **404**, 441 (1993).
- [72] LSST Science Collaboration *et al.*, [arXiv:0912.0201](https://arxiv.org/abs/0912.0201).
- [73] C. Chang, M. Jarvis, B. Jain, S. M. Kahn, D. Kirkby, A. Connolly, S. Krughoff, E. Peng, and J. R. Peterson, *Mon. Not. R. Astron. Soc.* **434**, 2121 (2013).
- [74] B. D. Sherwin and M. Schmittfull, *Phys. Rev. D* **92**, 043005 (2015).
- [75] B. Yu, J. Hill, and B. D. Sherwin, *Phys. Rev. D* **96**, 123511 (2017).
- [76] C. M. Hirata and U. Seljak, *Phys. Rev. D* **70**, 063526 (2004).
- [77] S. Bridle and L. King, *New J. Phys.* **9**, 444 (2007).
- [78] M. Brown, A. Taylor, N. Hambly, and S. Dye, *Mon. Not. R. Astron. Soc.* **333**, 501 (2002).
- [79] S. Samuroff *et al.* (DES Collaboration), *Mon. Not. R. Astron. Soc.* **489**, 5453 (2019).
- [80] H. Johnston *et al.*, *Astron. Astrophys.* **624**, A30 (2019).
- [81] J. Yao, H. Shan, P. Zhang, J.-P. Kneib, and E. Jullo, *Astrophys. J.* **904**, 135 (2020).
- [82] R. Mandelbaum *et al.* (LSST Dark Energy Science Collaboration), [arXiv:1903.09323](https://arxiv.org/abs/1903.09323).
- [83] Z. Vlah, N. E. Chisari, and F. Schmidt, *J. Cosmol. Astropart. Phys.* **01** (2020) 025.

- [84] M. C. Fortuna, H. Hoekstra, B. Joachimi, H. Johnston, N. E. Chisari, C. Georgiou, and C. Mahony, *Mon. Not. R. Astron. Soc.* **501**, 2983 (2021).
- [85] M. Lewandowski, A. Perko, and L. Senatore, *J. Cosmol. Astropart. Phys.* **05** (2015) 019.
- [86] S.-F. Chen, E. Castorina, and M. White, *J. Cosmol. Astropart. Phys.* **06** (2019) 006.
- [87] E. Semboloni, H. Hoekstra, and J. Schaye, *Mon. Not. R. Astron. Soc.* **434**, 148 (2013).
- [88] A. Mead, T. Tröster, C. Heymans, L. Van Waerbeke, and I. McCarthy, *Astron. Astrophys.* **641**, A130 (2020).
- [89] A. Mead, S. Brieden, T. Tröster, and C. Heymans, [arXiv:2009.01858](https://arxiv.org/abs/2009.01858).
- [90] S. N. Debackere, J. Schaye, and H. Hoekstra, *Mon. Not. R. Astron. Soc.* **492**, 2285 (2020).
- [91] T. Eifler, E. Krause, S. Dodelson, A. Zentner, A. Hearin, and N. Gnedin, *Mon. Not. R. Astron. Soc.* **454**, 2451 (2015).
- [92] I. Mohammed and N. Y. Gnedin, *Astrophys. J.* **863**, 173 (2018).
- [93] H.-J. Huang, T. Eifler, R. Mandelbaum, and S. Dodelson, *Mon. Not. R. Astron. Soc.* **488**, 1652 (2019).
- [94] A. Schneider and R. Teyssier, *J. Cosmol. Astropart. Phys.* **12** (2015) 049.
- [95] 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.
- [96] G. Aricò, R. E. Angulo, C. Hernández-Monteagudo, S. Contreras, M. Zennaro, M. Pellejero-Ibañez, and Y. Rosas-Guevara, *Mon. Not. R. Astron. Soc.* **495**, 4800 (2020).
- [97] G. Aricò, R. E. Angulo, C. Hernández-Monteagudo, S. Contreras, and M. Zennaro, [arXiv:2009.14225](https://arxiv.org/abs/2009.14225).
- [98] A. Schneider, N. Stotra, A. Refregier, A. J. Weiss, M. Knabenhans, J. Stadel, and R. Teyssier, *J. Cosmol. Astropart. Phys.* **04** (2020) 019.
- [99] T. Tröster, C. Ferguson, J. Harnois-Déraps, and I. G. McCarthy, *Mon. Not. R. Astron. Soc.* **487**, L24 (2019).
- [100] F. Villaescusa-Navarro *et al.*, [arXiv:2010.00619](https://arxiv.org/abs/2010.00619).
- [101] A. Cooray and R. K. Sheth, *Phys. Rep.* **372**, 1 (2002).
- [102] G. Parimbelli, M. Viel, and E. Sefusatti, *J. Cosmol. Astropart. Phys.* **01** (2019) 010.
- [103] A. H. Gonzalez, S. Sivanandam, A. I. Zabludoff, and D. Zaritsky, *Astrophys. J.* **778**, 14 (2013).
- [104] S. Grandis, J. J. Mohr, J. P. Dietrich, S. Bocquet, A. Saro, M. Klein, M. Paulus, and R. Capasso, *Mon. Not. R. Astron. Soc.* **488**, 2041 (2019).
- [105] B. Soergel *et al.* (DES and SPT Collaborations), *Mon. Not. R. Astron. Soc.* **461**, 3172 (2016).
- [106] V. Vikram, A. Lidz, and B. Jain, *Mon. Not. R. Astron. Soc.* **467**, 2315 (2017).
- [107] H. Tanimura, S. Zaroubi, and N. Aghanim, *Astron. Astrophys.* **645**, A112 (2021).
- [108] S. Lim, H. Mo, H. Wang, and X. Yang, *Astrophys. J.* **889**, 48 (2020).
- [109] N. Battaglia, S. Ferraro, E. Schaan, and D. Spergel, *J. Cosmol. Astropart. Phys.* **11** (2017) 040.
- [110] H. Tanimura, G. Hinshaw, I. G. McCarthy, L. Van Waerbeke, N. Aghanim, Y.-Z. Ma, A. Mead, T. Tröster, A. Hojjati, and B. Moraes, *Mon. Not. R. Astron. Soc.* **491**, 2318 (2020).
- [111] S. Amodeo *et al.*, *Phys. Rev. D* **103**, 063514 (2021).
- [112] Y.-Z. Ma, Y. Gong, T. Troster, and L. Van Waerbeke, *Mon. Not. R. Astron. Soc.* **500**, 1806 (2020).
- [113] S. Pandey, E. Baxter, and J. Hill, *Phys. Rev. D* **101**, 043525 (2020).
- [114] N. Battaglia *et al.*, [arXiv:1903.04647](https://arxiv.org/abs/1903.04647).
- [115] A. Hojjati, I. G. McCarthy, J. Harnois-Déraps, Y.-Z. Ma, L. Van Waerbeke, G. Hinshaw, and A. M. Le Brun, *J. Cosmol. Astropart. Phys.* **10** (2015) 047.
- [116] A. Hojjati *et al.*, *Mon. Not. R. Astron. Soc.* **471**, 1565 (2017).
- [117] J. Hill and D. N. Spergel, *J. Cosmol. Astropart. Phys.* **02** (2014) 030.