Inference of gravitational lensing and patchy reionization with future CMB data

Federico Bianchini[1,2,3,*] and Marius Millea[4,5,†]

[1]*Kavli Institute for Particle Astrophysics and Cosmology, Stanford University, 452 Lomita Mall, Stanford, California 94305, USA*
[2]*SLAC National Accelerator Laboratory, 2575 Sand Hill Road, Menlo Park, California 94025, USA*
[3]*Department of Physics, Stanford University, 382 Via Pueblo Mall, Stanford, California 94305, USA*
[4]*Department of Physics, University of California, Berkeley, California 94720, USA*
[5]*Department of Physics, University of California, Davis, California 95616, USA*

We develop an optimal Bayesian solution for jointly inferring secondary signals in the cosmic microwave background (CMB) originating from gravitational lensing and from patchy screening during the epoch of reionization. This method is able to extract full information content from the data, improving upon previously considered quadratic estimators for lensing and screening. We forecast constraints using the marginal unbiased score expansion method, and show that they are largely dominated by CMB polarization, and depend on the exact details of reionization. For models consistent with current data which produce the largest screening signals, a detection (3σ) of the cross-correlation between lensing and screening is possible with SPT-3G and a detection of the autocorrelation is possible with CMB-S4. Models with the lowest screening signals evade the sensitivity of SPT-3G but are still possible to detect with CMB-S4 via their lensing cross-correlation.

I. INTRODUCTION

The large-scale structure of the Universe is backlit by cosmic microwave background (CMB) photons as they travel from the last scattering surface towards us. Maps of the CMB anisotropies can therefore be used to image gravitational potentials—through weak gravitational lensing, integrated Sachs-Wolfe and Rees-Sciama effects—and the gas distribution—through Thomson and inverse Compton scattering processes like the Sunyaev-Zel'dovich effects (e.g., [1–5]). The current generation of CMB surveys such as *Planck* [6], South Pole Telescope (SPT) [7] and Atacama Cosmology Telescope (ACT) [8] has started to tap into the promising potential of these CMB secondary anisotropies. Next-generation experiments—including the Simons Observatory (SO) [9], FYST/CCAT-prime [10], and CMB-S4 [11]—will provide a transformative high fidelity view of the secondary CMB anisotropies in intensity and polarization over large areas of the sky, revealing fundamental insights into both cosmology and astrophysics [12].

Mapping out the spatial distribution of diffuse ionized gas throughout the Universe can, for example, help us understand the physics of reionization, at high redshifts ($z \gtrsim 6$), and of the intergalactic medium, at lower redshifts ($z \lesssim 6$) (e.g., [13,14]). One approach to achieve this is by searching for the characteristic spatially-dependent suppression of the CMB temperature and polarization anisotropies produced by different scattering histories along different lines of sight, an effect known as "patchy screening." The magnitude of the effect is proportional to $e^{-\tau(\hat{n})}$, where $\tau(\hat{n})$ is the direction-dependent optical depth.

The so-called "quadratic estimator" (QE) has become the workhorse for extracting sources of statistical anisotropies, such as patchy screening and lensing, from CMB maps (e.g., [15–17]). The QE in the context of inhomogeneous optical depth reconstruction has been introduced by Dvorkin and Smith [18] and applied to WMAP and *Planck* data in [19–21] but the spatial fluctuations of τ have not been detected yet. While the QE has been successfully used on current datasets, it has some shortcomings. First, the presence of other distorting fields, like lensing, point sources, and inhomogeneous noise, will introduce additional non-Gaussianities in the data which in turn, lead to biases in the reconstructed field [22]. "Bias-hardened" estimators offer a solution to this problem at the cost of a signal-to-noise (S/N) ratio degradation, which can be as large as $\approx 40\%$ (e.g., [21,22]). Second, the QE will become significantly suboptimal at the instrumental noise levels soon-to-be reached by the most sensitive experiments [23,24]. At these depths, secondary anisotropies, rather than instrumental noise, limit the variance of the estimated

*[*]federico.bianxini@gmail.com
*[†]mariusmillea@gmail.com

field. To improve upon the QE, a variety of methods based on the full CMB Bayesian posterior have been proposed to extract the higher-order information and restore near optimality (e.g., [25–30]). Machine-learning approaches are also being currently investigated, but while promising, additional work towards the characterization of these methods and their biases is needed before they can be reliably applied to real data [31,32].

In this paper, we develop a complete Bayesian solution that unifies the optimal inference of the optical depth τ and CMB lensing potential ϕ together with delensing and cosmological parameter inference. Our method presents a number of appealing features. First of all, by making use of the full Bayesian posterior, the method is capable of optimally extracting the information content in CMB data at all noise levels. Furthermore, by simultaneously forward modeling the effects of lensing and screening on CMB observables, we are able to naturally account for any contamination (of lensing to screening and vice versa) in the reconstruction and to enhance the sensitivity to screening through the cosmic variance reduction due to delensing. Our method can also internally measure the correlation $\langle \phi\tau \rangle$ between the CMB lensing potential and optical depth fluctuations, which is larger than $\langle \tau\tau \rangle$, contains additional information on the relation between the ionized gas and dark matter distribution, and is expected to be detected with upcoming CMB surveys [33]. Finally, the CMB data are effectively "descreened" by our procedure, which can in turn mitigate any residual B-mode bias from screening to tensor-to-scalar ratio r searches [although the contamination is expected to be at the level of $r \sim 10^{-4}$ for standard reionization histories, see, e.g., [34,35], therefore not a significant concern for experiments targeting $\mathcal{O}(10^{-3})$].

We begin the paper with a review of the theoretical background and the effect of patchy screening on CMB observables in Sec. II. We present our method and illustrate its performance in Sec. III before concluding in Sec. IV. In the Appendix we provide a more detailed description of the halo model used to predict the patchy screening signal.

II. MODELING

A. Optical depth

The electron scattering optical depth measures the integrated electron density along the line-of-sight and is given by

$$\tau(\hat{\mathbf{n}}) = \sigma_T \int d\chi \, a \, n_e(\hat{\mathbf{n}}, \chi), \tag{1}$$

where $n_e(\hat{\mathbf{n}}, \chi)$ is the free electron number density at comoving distance χ along the direction $\hat{\mathbf{n}}$, σ_T is the Thomson scattering cross section, and a is the scale factor. The mean number density of free electrons can be expressed as $\bar{n}_e = (1 - 3/4Y_P)\bar{\rho}_{b0}/m_p a^{-3}\bar{x}_e = \bar{n}_{p0}\bar{x}_e/a^3$, where Y_P is the primordial helium abundance, ρ_{b0} is the

present-day baryon density, the proton mass is m_p, $\bar{x}_e$ is the mean ionization fraction, and we assumed that helium is singly ionized. We model the evolution of the volume-averaged ionization fraction $\bar{x}_e$ using a simple tanh fitting function parametrized by the redshift of reionization $z_{\rm re}$, defined as the redshift at which $\bar{x}_e$ is half of its maximum, and the duration of reionization Δ_z, i.e., the difference between the redshifts at which the Universe is 5% and 95% reionized, [36]

$$\bar{x}_e(z) = \frac{1}{2}\left[1 + \tanh\left(\frac{y_{\rm re} - y}{\Delta y_{\rm re}}\right)\right], \tag{2}$$

where $y(z) = (1 + z)^{3/2}$, $y_{\rm re} = y(z_{\rm re})$, and $\Delta y_{\rm re} = \sqrt{1 + z_{\rm re}}\Delta_z$.

Perturbations in the free electron number density δn_e can be sourced both by ionization fluctuations $\delta_x = \delta x_e/\bar{x}_e$ and by inhomogeneities in the gas density δ [14]. The former are only generated during the epoch of reionization, while the latter are produced both when reionization occurs, as well as in the postreionization universe. Fluctuations in the free electron density will then induce spatial fluctuations in the optical depth (e.g., [37]),

$$\tau(\hat{\mathbf{n}}) = \sigma_T n_{p0} \int \frac{d\chi}{a^2}\bar{x}_e(\chi)(1 + \delta_x(\hat{\mathbf{n}}, \chi))(1 + \delta(\hat{\mathbf{n}}, \chi)) \tag{3}$$

$$= \bar{\tau} + \delta\tau(\hat{\mathbf{n}}). \tag{4}$$

The contribution to the τ anisotropies from the spatial distribution of free electrons in galaxies and clusters is then given by the redshifts below in which $\bar{x}_e = 1$.

Under the Limber approximation [38], valid for angular scales $\ell \gg 20$ relevant for this work, the angular power spectrum of the optical depth fluctuations can be evaluated as

$$C_\ell^{\tau\tau} = \sigma_T^2 n_{p0}^2 \int \frac{d\chi}{a^4\chi^2} P_{\delta_e\delta_e}(k_\ell, \chi), \tag{5}$$

where $P_{\delta_e\delta_e}$ is the power spectrum of the density-weighted ionization fraction fluctuations, and $k_\ell = (\ell + \frac{1}{2})/\chi$.

The quantity $P_{\delta_e\delta_e}$ encodes all the relevant astrophysical aspects of the epoch of reionization, including its morphology and timing. Modeling reionization and the spatial distribution of free electrons is a challenging task due to the complexity of the physical processes involved and the limited observational access we currently have to those cosmic epochs. Therefore, instead of specifying a given physical mechanism of reionization (e.g., [39]), we choose to adopt the "bubble model" introduced in [14,18,40] that allows us to phenomenologically parametrize the HII spectrum, $P_{\delta_e\delta_e}$. In this framework, the HII regions around the ionizing sources, such as galaxies or quasars, are biased tracers of the underlying dark matter halos. The ionized bubbles are assumed to be spherical with an average radius $\bar{R}$ (in Mpc), while their radii distribution is modeled as a

log-normal distribution of width $\sigma_{\ln R}$, i.e., skewed towards smaller bubble sizes. The size and evolution of these HII bubbles are sensitive to the mass and brightness of the ionizing sources. As time advances, the HII bubbles grow in size and percolate, eventually leading to a complete reionization of the intergalactic medium. In the Appendix we provide a more detailed discussion of the halo model.

To summarize, our inhomogeneous reionization modeling is described by a set of four parameters: the redshift and duration of reionization $\{z_{\rm re}, \Delta_z\}$, which specify the mean ionization history $\bar{x}_e(z)$, and the characteristic size and standard deviation of the log-normal bubble radius distribution, $\{\bar{R}, \sigma_{\ln R}\}$.

B. Cross-correlation with CMB lensing

In this work we are also interested in evaluating the correlation between the optical depth fluctuations and the integrated matter distribution along the line-of-sight as traced by the CMB lensing potential $\phi(\hat{\mathbf{n}})$ defined as (e.g., [3])

$$\phi(\hat{\mathbf{n}}) = -2 \int d\chi \frac{\chi - \chi_*}{\chi \chi_*} \Psi(\hat{\mathbf{n}}, \chi). \tag{6}$$

Here, $\Psi(\hat{\mathbf{n}}, \chi)$ is the Weyl (gravitational) potential, that in standard cosmologies can be directly related to the comoving matter perturbations $\delta(\hat{\mathbf{n}}, \chi)$ through the Poisson equation, and χ_* is the comoving distance to the last scattering surface at $z_* \simeq 1090$.

A cross-correlation between optical depth fluctuations and CMB lensing is naturally expected since the same dark matter halos signposted by the HII bubbles at high redshift ($z \gtrsim 6$) and by galaxies and galaxy clusters at low redshift ($z \lesssim 6$) act as lenses for the CMB photons. In fact, the CMB lensing kernel has nonzero support out to high redshift and it is about half of its maximum around $z \approx 10$.

Using Eq. (6), the cross-power spectrum between the fluctuating part of τ and CMB lensing potential takes the following form:

$$C_\ell^{\phi\tau} = \frac{3H_0^2 \Omega_m \sigma_T n_{p0}}{c\ell^2} \int \frac{d\chi}{a^3} \frac{\chi - \chi_*}{\chi \chi_*} P_{\delta\delta_e}(k_\ell, \chi), \tag{7}$$

where H_0 is the Hubble constant, and $P_{\delta\delta_e}(k_\ell, \chi)$ denotes the three-dimensional cross spectrum between the matter density contrast δ and the free electron fluctuations δ_e.

C. Fiducial scenarios

While we do not specify a given physical model for reionization, here we consider two limiting reionization scenarios to study the performance of our approach and to get a sense for the expected S/N ratio with current and upcoming surveys.

In the first case (optimistic model), reionization starts early on, proceeds for an extended period of time, and is driven by larger bubbles with a broad radii distribution

spread. For this scenario we fix $\{z_{\rm re}, \Delta_z, \bar{R}, \sigma_{\ln R}\} = \{10, 4, 5, \ln(2)\}$. Aside from giving a signal which is large but not currently detectable, a value of $\sigma_{\ln R} = \ln(2)$, is on the upper end of allowed values given existing cross-correlations between QE reconstructions of τ and Compton y-maps using *Planck* data [21]. In the second case (pessimistic model), reionization occurs at lower redshift, and its duration is shorter, while the bubbles that reionize the Universe are smaller in size, and their radii distribution is narrower. When analyzing this scenario we set $\{z_{\rm re}, \Delta_z, \bar{R}, \sigma_{\ln R}\} = \{7, 1, 1, 0.1\}$. As we will see, the nomenclature of optimistic or pessimistic refers to the prospects of detecting the patchy screening effect in the two cases.

In the left-hand panel of Fig. 1, we show the two reionization histories. The higher redshift of reionization in the optimistic model results in a larger integrated mean optical depth $\bar{\tau}$ than in the pessimistic model. Both values are consistent within 2σ with the constraint from *Planck* on $\bar{\tau} = 0.058 \pm 0.012$ [6]. The middle and right panels of Fig. 1 show the optical depth fluctuation angular autospectrum and its cross-correlation with CMB lensing, respectively. Qualitatively, a longer duration of reionization translates to more pronounced fluctuations in τ, corresponding to a larger overall amplitude of $C_\ell^{\tau\tau}$ and $C_\ell^{\phi\tau}$. The peak of these spectra is mostly determined by the effective radius of the HII bubbles, with larger bubbles enhancing the power spectra at lower multipoles.

D. Hierarchical model for CMB data

The patchy nature of the reionization process at $z \gtrsim 6$ and the clumpiness of the spatial distribution of free electrons at lower redshifts ($z \lesssim 6$) lead to an anisotropic optical depth, $\tau = \tau(\hat{\mathbf{n}})$. In turn, spatial variations of the optical depth will introduce new observational signatures in the CMB sky.

Three main effects can be identified. First, Thomson scattering of the remote CMB temperature quadrupole by ionized bubbles generates new polarization (e.g., [13,14,41–44]).[1] Second, the radial peculiar velocity of ionized bubbles relative to the observer sources a temperature fluctuation through the kinematic Sunyaev-Zel'dovich (KSZ) effect (e.g., [41,44–48]). Third, anisotropies in the optical depth will induce a spatially-dependent screening of the primary CMB anisotropies caused by the CMB photons being scattered into and out of our line-of-sight (e.g., [43,49]). In this paper, we focus exclusively on the latter effect.

Considering an unperturbed CMB field $f \in \{T, Q, U\}$, the map-level effect of anisotropic screening is to spatially modulate the amplitude of the CMB temperature and polarization anisotropies by a factor $e^{-\tau(\hat{\mathbf{n}})}$ as

---

[1]While new E-mode polarization is generated even in the case of homogeneous reionization, additional B-modes can only be sourced if the distribution of free electrons is inhomogeneous.
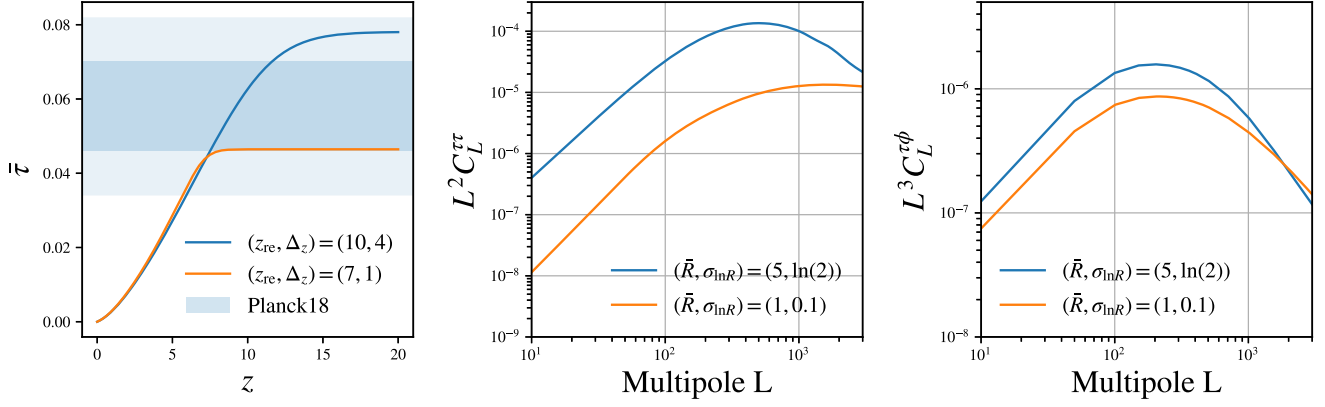
FIG. 1. Left: redshift evolution of the electron scattering optical depth based on the ionization fraction model of Eq. (2) assuming two different reionization histories (blue and orange solid lines). The shaded blue bands show the 1 and 2σ constraints on $\bar{\tau}$ from *Planck* [6]. Middle: angular power spectrum of the optical depth calculated from the bubble model. The colors refer to the reionization history parameters of the left panel. Right: cross-power spectrum between optical depth and CMB lensing potential (same color coding as the middle panel).

$$(\mathbb{S}(\tau)f)(\hat{\mathbf{n}}) = e^{-\tau(\hat{\mathbf{n}})}f(\hat{\mathbf{n}}), \tag{8}$$

where we have introduced the screening operator $\mathbb{S}(\tau)$.

On the other hand, the gravitational potentials associated to the intervening matter distribution between us and the last scattering surface induce a remapping of the primary CMB anisotropies, which we model (in the absence of screening) as

$$(\mathbb{L}(\phi)f)(\hat{\mathbf{n}}) = f(\hat{\mathbf{n}} + \nabla\phi(\hat{\mathbf{n}})), \tag{9}$$

where we have similarly introduced the lensing operator $\mathbb{L}(\phi)$.

In our treatment of the two effects, we make the simplifying assumption that patchy screening happens first, and the lensing operator is then applied to the screened CMB fields,

$$(\mathbb{L}(\phi)\mathbb{S}(\tau)f)(\hat{\mathbf{n}}) = e^{-\tau(\hat{\mathbf{n}}+\nabla\phi(\hat{\mathbf{n}}))}f(\hat{\mathbf{n}} + \nabla\phi(\hat{\mathbf{n}})). \tag{10}$$

Given these operators, we can model observed CMB data as

$$d = \mathbb{A}\mathbb{L}(\phi)\mathbb{S}(\tau)f + n, \tag{11}$$

where d is the data, f are the unlensed and unscreened CMB fields, and n is the instrumental noise. The linear operator $\mathbb{A} \equiv \mathbb{M}\mathbb{T}\mathbb{B}$ encodes the effects of the instrumental beam $\mathbb{B}$, the transfer function $\mathbb{T}$, and any pixel mask $\mathbb{M}$.

The goal of this analysis is to simultaneously extract lensing and patchy screening effects from noisy CMB data. There are different ways to cast this problem. In this work we choose to introduce three amplitude parameters $\theta = \{A_{\phi\phi}, A_{\tau\tau}, A_{\phi\tau}\}$ that rescale some fiducial CMB lensing and optical depth fiducial isotropic power

spectra as $C_\ell^{XY} = A_{XY}C_\ell^{XY,\mathrm{fid}}$ with $X, Y \in \{\phi, \tau\}$. This allows us to rapidly explore the reionization parameter space and to efficiently forecast the detection capabilities of various surveys in different scenarios. While we adopt this approach for the scope of this paper, we note that nothing prevents us from carrying out the statistical inference directly at the power spectrum level (e.g., [30]), or on the phenomenological parameters $(\bar{R}, \sigma_{\ln R})$ directly.

With this definition, the CMB screening and lensing problem can then be formulated as

$$(A_{\phi\phi}, A_{\tau\tau}, A_{\phi\tau}) \sim \mathcal{U}(0, \infty), \tag{12}$$

$$f \sim \mathcal{N}(0, \mathbb{C}_f), \tag{13}$$

$$\phi, \tau \sim \mathcal{N}(0, \mathbb{C}_{\phi\tau}(\theta)), \tag{14}$$

$$d \sim \mathcal{N}(\mathbb{L}(\phi)\mathbb{S}(\tau)f, \mathbb{C}_n), \tag{15}$$

where $\mathcal{U}$ and $\mathcal{N}$ denote uniform and Gaussian distributions, respectively. We assume that both the unperturbed CMB and the experimental noise are Gaussian random fields described by covariance operators $\mathbb{C}_n$ and $\mathbb{C}_f$, respectively, where the latter reads

$$\mathbb{C}_f = \begin{bmatrix} C_\ell^{TT} & C_\ell^{TE} & 0 \\ C_\ell^{TE} & C_\ell^{EE} & 0 \\ 0 & 0 & C_\ell^{BB} \end{bmatrix}. \tag{16}$$

Here we adopt the prior that both ϕ and τ are correlated Gaussian random fields with covariance given by

TABLE I. Technical specifications for the CMB surveys considered in this work. We show the frequency band centers (in GHz), full width at half maximum (FWHM) apertures (in arc min), temperature white-noise levels $\Delta_T[\mu\text{K-arc min}]$ (the polarization white-noise level is $\Delta_P = \sqrt{2}\Delta_T$), and atmospheric noise parameters (middle and bottom rows in each subpanel show the numbers for temperature and polarization, respectively).

Experiment	f_{sky}	Beam θ_{FWHM} [arc min]						Noise $\Delta_T[\mu\text{K-arc min}]$					
		30	40	90	150	220	270	30	40	90	150	220	270
SPT-3G	3.6%			1.7	1.2	1.0				3.0	2.2	8.8	
										(1200,3.0)	(1200,4.0)	(1200,4.0)	
										(300,1)			
SO	40%	7.4	5.1	2.2	1.4	1.0	0.9	71.0	36.0	8.0	10.0	22.0	54.0
										(1000,3.5)			
										(700,1.4)			
S4-wide	67%	7.3	5.5	2.3	1.5	1.0	0.8	21.8	12.4	2.0	2.0	6.9	16.7
								(471,3.5)	(428,3.5)	(2154,3.5)	(4364,3.5)	(7334,3.5)	(7308,3.5)
										(700,1.4)			
S4-deep	3%	8.4	5.8	2.5	1.6	1.1	1.0	4.6	2.94	0.45	0.41	1.29	3.07
								(1200,4.2)	(1200,4.2)	(1200,4.2)	(1900,4.1)	(2100,3.9)	(2100,3.9)
								(150,2.7)	(150,2.7)	(150,2.6)	(200,2.2)	(200,2.2)	(200,2.2)

$$\mathbb{C}_{\phi\tau}(\theta) = \begin{bmatrix} A_{\phi\phi}C_\ell^{\phi\phi} & A_{\phi\tau}C_\ell^{\phi\tau} \\ A_{\phi\tau}C_\ell^{\phi\tau} & A_{\tau\tau}C_\ell^{\tau\tau} \end{bmatrix}, \quad (17)$$

where we emphasize the explicit dependence of $\mathbb{C}_{\phi\tau}$ on the spectral amplitude parameters θ. Although the τ field is not expected to be Gaussian, due to the nonlinear growth of the ionizing bubbles, our choice to take a Gaussian prior is made for simplicity and because it is in some sense the conservative choice for a first detection. In this way, inference will not attempt to extract any information from higher-order statistics, which are more difficult to realistically model.

With these ingredients, the joint likelihood distribution takes the following form:

$$\mathcal{P}(d, \underbrace{f, \phi, \tau}_{z} | \underbrace{A_{\tau\tau}, A_{\phi\tau}, A_{\phi\phi}}_{\theta})$$

$$\propto \frac{\exp\left\{-\frac{(d-\mathbb{AL}(\phi)\mathbb{S}(\tau)f)^2}{2\mathbb{C}_n}\right\}}{\det \mathbb{C}_n^{1/2}} \frac{\exp\left\{-\frac{f^2}{2\mathbb{C}_f}\right\}}{\det \mathbb{C}_f^{1/2}} \frac{\exp\left\{-\frac{(\phi\oplus\tau)^2}{2\mathbb{C}_{\phi\tau}(\theta)}\right\}}{\det \mathbb{C}_{\phi\tau}(\theta)^{1/2}},$$

$$(18)$$

where we have explicitly highlighted that the latent space z, over which marginalization will be performed, is composed by $\{f, \phi, \tau\}$, and we recall that the parameters of interest θ are the power spectra amplitudes $\{A_{\phi\phi}, A_{\tau\tau}, A_{\phi\tau}\}$. In Eq. (18), we denote the concatenation of the two abstract vectors ϕ and τ using the symbol $\oplus$, and we use the shorthand $x^2/\mathbb{N} \equiv x^\dagger \mathbb{N}^{-1}x$.

We give the posterior function in Eq. (18) for exposition, but we note that our code, implemented in

CMBLensing.jl,[2] calculates it automatically using a probabilistic programming language given only the forward model in Eq. (12), and calculates posterior gradients using automatic differentiation [29]. This highlights the strength of the Bayesian approach, which allows forward models to be quickly turned into pipelines ready to use for inference.

III. RESULTS

A. Experimental setups

We consider a number of CMB surveys representative of the current and upcoming experimental landscape. The experimental specifications are listed in Table I (taken from [50,51]).

Our noise spectra modeling includes a contribution from both instrumental and atmospheric noise as

$$N_\ell = \Delta_T^2\left[1 + \left(\frac{\ell}{\ell_{\text{knee}}}\right)^{-\alpha_{\text{knee}}}\right], \quad (19)$$

where Δ_T^2 is the detector white noise, while ℓ_{knee} and α_{knee} parametrize the atmospheric $1/f$ noise.

We do not include emission from astrophysical foregrounds in our simulations but incorporate their impact by considering the effective noise levels after component separation through harmonic-space internal linear combination (ILC) methods (e.g., [52,53]). We model the total covariance between different frequency channels at each multipole ℓ as a linear mixture of CMB, foregrounds, and noise contributions,

---

[2]https://github.com/marius311/CMBLensing.jl

$$\mathbf{C}_\ell = \mathbf{e}\mathbf{e}^\dagger C_\ell^{\mathrm{CMB}} + \mathbf{C}_\ell^{\mathrm{FG}} + \mathbf{N}_\ell, \qquad (20)$$

where $\mathbf{e}$ is a column vector of all ones (for maps in thermodynamic temperature units). The foreground contributions in intensity include emission from radio point sources, the cosmic infrared background, and diffuse kinematic and thermal Sunyaev-Zel'dovich signals. The polarized CMB sky contains emissions from galactic dust and synchrotron. For simplicity, we adopt the "standard" ILC noise that minimizes the total variance of the final CMB + KSZ map, so that the postcomponent separation ILC noise curve can be calculated as follows:

$$N_\ell^{\mathrm{ILC}} = \left[\sum_{ij}(\mathbf{C}^{-1})_\ell^{ij}\right]^{-1}. \qquad (21)$$

Our forecasts are performed on 512×512 pixel flat-sky maps with 2 arc min pixels, covering roughly 300 deg^2 of sky. These are then rescaled by an f_{sky} factor to match the expected sky coverage of each survey. Because both lensing and patchy screening are spatially local operations, we expect this scaling will give results very similar to having performed the forecasts on larger sky areas directly, but at much lower computational cost. We do not include any masking, which we expect will be a <10% effect [30].

B. Joint MAPs

As a first check, we compute the joint maximum *a posteriori* (MAP) estimate of the gravitational potential ϕ and the optical depth τ on a suite of simulated data. MAP estimates are only an intermediate step towards forecasting or inference since they have biased power spectra, but they are useful as a sanity check and to build intuition about the problem. They are computed by fixing the spectral amplitudes to some fiducial value (here unity) and maximizing the posterior probability,

$$(\hat{f}_J, \hat{\phi}_J, \hat{\tau}_J) = \underset{f,\phi,\tau}{\mathrm{argmax}} \, \log \mathcal{P}(f,\phi,\tau|\theta,d). \qquad (22)$$

We iteratively maximize the posterior using the coordinate descent algorithm introduced in [28], with a modified Newton-Raphson step that performs joint (ϕ, τ) updates.

Figure 2 shows a typical $\hat{\tau}_J$ as compared to the simulation truth, assuming S4-deep noise levels and using only polarization. The contour levels are set at 0 and ± 1 times the pixel root mean square of the MAP reconstruction and can be visibly seen to trace the true field, with regions of low τ encircled by blue contour regions of high τ encircled by orange contours. More quantitatively, in Fig. 3, we show the cross-correlation coefficient ρ_ℓ at different scales between the true (ϕ, τ) maps and the corresponding joint MAP estimates from the simulated data, averaged over several realizations. We see that the MAP ϕ reconstruction is correlated to the true ϕ map, as expected, and that the MAP τ reconstruction is correlated to the true τ, indicating
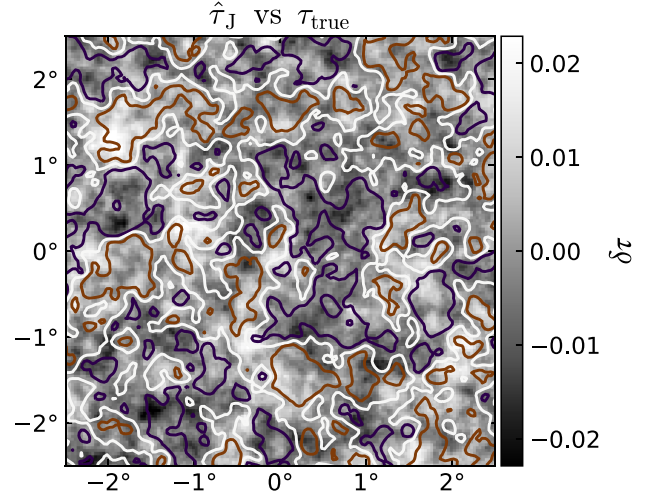

FIG. 2. An example of MAP optical depth reconstruction $\hat{\tau}_J$ shown as contours, compared to the true field for the input simulation τ_{true} shown as the gray scale map. We assume S4-deep noise levels and here, for clarity, show just a small cutout of the larger maps, which are reconstructed as part of the forecasting procedure. Some partial correlation is visible by eye.

the potential for a detection. We also see that the MAP τ is correlated to the true ϕ. While this is also expected to some degree, given that the true τ and ϕ fields are intrinsically correlated, the fact that $\langle\tau^{\mathrm{true}}\phi^{\mathrm{MAP}}\rangle \neq \langle\tau^{\mathrm{MAP}}\phi^{\mathrm{true}}\rangle$ implies an additional lensing-induced bias in the power spectra of the τ MAPs (beyond the suppression of power which is typical of MAP estimates). This is not surprising, as it is known that the τ QE contains a lensing-induced bias [21,22], and MAP estimates asymptote to the QE in the high-noise limit. We note that our Bayesian approach implicitly deals with this when marginalizing over the $\{f, \phi, \tau\}$ maps, allowing unbiased inference of parameters, which control the theory auto and cross-correlation spectra without having to explicitly compute any such bias. As an additional check, in Fig. 3 we include the correlation coefficients between the true τ and MAP ϕ (and vice versa) when we turn off the intrinsic correlation between the true τ and ϕ fields.

C. Forecasting results

The key challenge for inference or forecasting is the marginalization over the $\{f, \phi, \tau\}$ fields. While we could use an exact sampling algorithm like Hamiltonian Monte Carlo (HMC) to perform this marginalization, we instead use the approximate but much faster marginal unbiased score expansion (MUSE) method [30,54,55]. For data vectors as large as ours here, MUSE is unbiased, but in theory can be suboptimal. However, it was demonstrated in [30] that for gravitational lensing inference, it achieves uncertainties which are <10% away from the optimal Fisher limit. We do not expect the addition of patchy screening to qualitatively change this behavior,
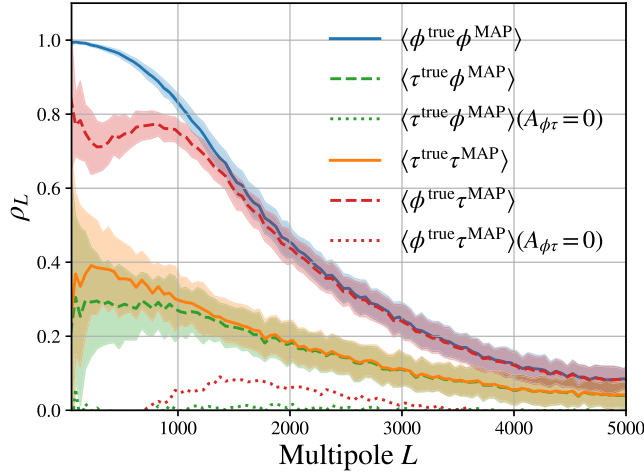
FIG. 3. Cross-correlation coefficient between the input CMB lensing ϕ and optical depth τ fields and the corresponding MAP estimates from simulated data, assuming S4-deep noise levels. The solid line represents the mean correlation coefficient across $N_{\mathrm{sim}} = 100$ simulations, while the shaded areas denote the 1σ scatter for a single realization. The average cross-correlation coefficient between the reconstructed MAP τ and input ϕ maps is shown as a dashed red line (the opposite case is denoted by the dashed green line). For reference, we also include the correlation coefficient between the reconstructed MAP and true maps when the input τ and ϕ fields are not intrinsically correlated, i.e., $A_{\phi\tau} = 0$, as the dotted lines (same color coding as the $A_{\phi\tau} = 1$ case).

motivating its use here. Although we expect our results are close to optimal, they should be regarded as an upper bound on possible uncertainties, and whether tighter constraints are possible could be checked by analyzing our posterior probability function with HMC.

Following [30], we use MUSE to compute the J and H matrices for the parameters set $\{A_{\phi\phi}, A_{\tau\tau}, A_{\phi\tau}\}$, then use these to form the marginal posterior covariance matrix. The J matrix is the covariance of gradients of the log posterior with respect to parameters at the joint MAP, and the H matrix is akin to a "response function," which quantifies how the gradient with respect to each parameter responds to injecting a shift in another parameter into the simulated data. In practice, the computation involves only obtaining MAP estimates of the kind pictured in Fig. 2 for a suite of data simulations, and evaluating gradients of the posterior in Eq. (18) at these points in parameter space. The MUSE covariance, much like the inverse Fisher matrix, is already averaged over data realizations and so directly gives us our desired forecast.

We begin by examining the information content in temperature vs polarization. In Fig. 4, we plot the forecasted covariance for the SPT-3G configuration with $\ell_{\mathrm{max}} = 3000$ as one- and two-dimensional marginalized contours. We find that there is almost no weight from temperature on constraints to the overall lensing or
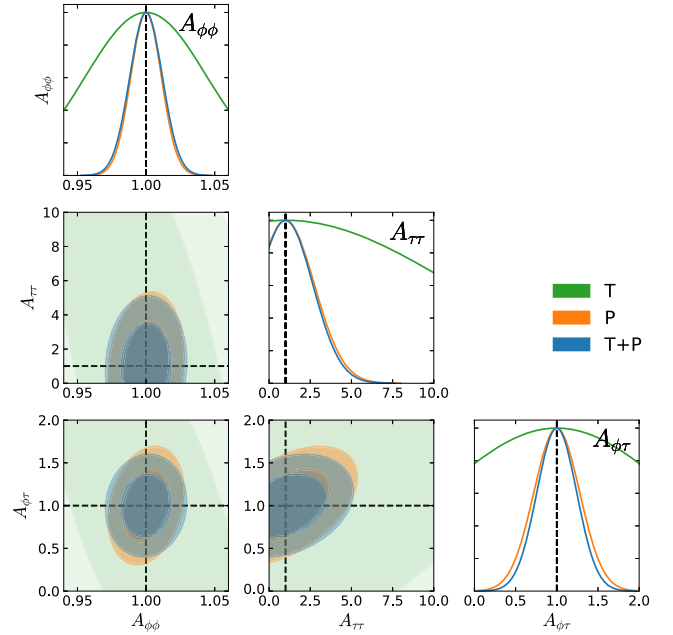

FIG. 4. Forecasted constraints on lensing and screening amplitudes from temperature and/or polarization for SPT-3G, assuming $\ell_{\mathrm{max}} = 3000$. We find that at these noise levels (and below) and temperature adds negligible information in terms of the overall amplitude parameters.

screening amplitudes. It is already known that at SPT-3G noise levels, the total lensing S/N ratio is highly dominated by polarization, and we find the same is true for the screening effect. We expect this will only get more
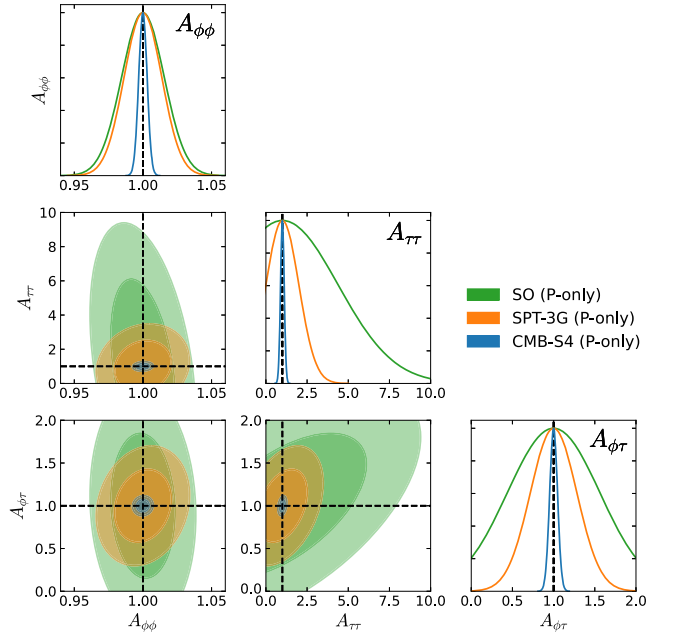

FIG. 5. Forecasted constraints from polarization for different experiments, given the optimistic reionization model and assuming $\ell_{\mathrm{max}} = 5000$.

TABLE II. Forecasted 1σ uncertainties from polarization for various configurations. Each entry gives a range corresponding to pessimistic and optimistic models (except in the SO case where only the optimistic model is considered). In all cases, the amplitude parameters are defined such that their fiducial value is unity (and note that the first column is multiplied by 100).

	$\ell_{\max}$	$100\,\sigma(A_{\phi\phi})$	$\sigma(A_{\tau\tau})$	$\sigma(A_{\phi\tau})$
SO	3000	1.5–	2.3–	0.51–
	5000	1.5–	2.2–	0.51–
SPT-3G	3000	1.4–1.4	0.97–3	0.29–0.44
	5000	1.4–1.4	0.95–3	0.28–0.43
S4-wide	3000	0.38–0.38	0.29–2.6	0.080–0.18
	5000	0.37–0.41	0.29–2.6	0.078–0.16
S4-deep	3000	0.57–0.55	0.13–1.3	0.066–0.110
	5000	0.55–0.51	0.12–1.3	0.060–0.088
S4	3000	0.31–0.31	0.12–1.2	0.050–0.088
	5000	0.31–0.30	0.11–1.2	0.047–0.072

drastic as noise levels decrease, thus for the forecasts described below, we consider only polarization for simplicity. By doing so, however, we are underestimating constraints possible with SO, which has higher noise levels and could receive a non-negligible contribution from temperature. Doing such a forecast would depend on more complex assumptions about temperature foregrounds, so we leave this to a future work.

We now consider the potential for detecting the screening signal from different ongoing and future CMB surveys. These results are summarized in Table II and Fig. 5. Below, we will use the term "detection" to refer to constraining a nonzero amplitude parameter at 3σ. With SPT-3G, it will be possible to make a detection of the patchy screening correlation with the lensing potential given the optimistic model, but not given the pessimistic model. With SO polarization, it is not possible to detect either model. In general, detecting the autocorrelation of the patchy screening is comparatively harder than detecting the cross spectrum and is not possible with SPT-3G or SO, regardless of model. For CMB-S4, either scenario leads to a guaranteed detection of the cross-correlation, and additionally, the autocorrelation is detectable given the optimistic model. We note that for CMB-S4, the deep survey contributes slightly more weight than the wide survey. With the pessimistic model, a guaranteed detection of the autocorrelation of the screening effect will require at least a fifth generation CMB survey.

D. Comparisons with other works

Forecasts similar to ours have been computed in previous works, based on the QE instead of on Bayesian inference. We do not attempt a careful quantitative comparison with other works but comment briefly on similarities.

Roy *et al.* [56] computed the total S/N ratio in the QE τ autospectrum, while arguing that lensing-induced biases were ignorable for their purposes due to sufficient S/N ratio. One of their forecasts considers a reionization model with $\bar{R} = 5$ and $\sigma_{\ln R} = \ln(2)$, similar to our optimistic model, with a similar total optical depth and lower assumed noise levels for S4. For this case, they predict a 5σ detection is possible with S4. Our S4 optimistic forecast gives a 10σ detection of $A_{\tau\tau}$, despite our higher noise assumptions. Feng and Holder [33] perform QE forecasts, which account for lensing-induced biases, and additionally consider the cross spectrum between lensing and screening fields. We have computed theory spectra using the reionization parameters quoted in their paper, visually verifying similar $C_\ell^{\phi\tau}$ spectra, and performed a forecast for S4 with their noise assumptions (more optimistic than ours). Their result gives a total S/N ratio of 22σ in the $\phi(EB)\tau(EB)$ estimator, and we find a 23σ detection of nonzero $A_{\phi\tau}$ for this case.

We also note that amplitude parameters only capture the total S/N ratio, and for spectra, which are very red at small scales such as the lensing and patchy screening spectra, the S/N ratio is mainly limited by the large-scale sample variance, not the reconstruction noise. This is the reason why we do not see more improvement with $\ell_{\max}$ in Table II and is also a reason why the above comparison is an incomplete test of the exact improvements upon the QE which are happening. An internal scale-dependent comparison of Bayesian and QE estimates of patchy screening, however, is outside the scope of this work, as forming an unbiased QE patchy reconstruction pipeline akin to [33] is significantly nontrivial, and regardless, we expect theoretically that the Bayesian pipeline will extract the same or more information than the QE pipeline. Our rough comparisons above indeed appear to be consistent with this expectation.

IV. CONCLUSIONS

In this work, we have developed a Bayesian model which can be used to simultaneously extract fluctuations in the optical depth τ and gravitational lensing potential ϕ together with delensed and descreened CMB fields, as well as the theory spectra which describe their statistics. Here, we have used MUSE to infer the (amplitudes of) the theory spectra and to forecast their uncertainties from upcoming probes. In the future, this code can directly be used to directly obtain estimates of the amplitudes from data, or can be used with HMC sampling or other inference tools to obtain joint inference on the spectra and fields together.

The Bayesian approach allows us to optimally recover all available information from the data and implicitly handles any bias subtraction, such as the lensing-induced bias of the τ QE. In practice, we find its strength is its simplicity, requiring us only to code up a field-level forward model,

which includes lensing and screening, with the posterior probability function and its gradients automatically derived, and MUSE providing quick forecasting, which requires only solving optimization problems. This type of analysis can readily be extended to include other secondary signals or foregrounds, and its equivalent could be applied to other problems in cosmology which have computable forward models. The trade-off is that any real analysis requires high-fidelity modeling of observed data, but these types of pipelines are being developed, including for the next generation of CMB experiments.

We have found that at the low-noise levels characteristic of the current most sensitive and upcoming CMB surveys ($\Delta_P \lesssim 5 \ \mu$K-arc min), polarization data carry the majority of the constraining power on patchy reionization. Owing to the faintness of the patchy screening effect, a definitive detection of the optical depth fluctuations will prove challenging in autocorrelation. Depending on the details of the reionization history, a next-generation experiment like CMB-S4 might be able to detect the signal, but for more pessimistic scenarios, we would require a futuristic Stage-5 mission. On the other hand, measuring the cross-correlation between ϕ and τ holds great promise for detecting optical depth anisotropies. For the full-depth SPT-3G survey, we forecast a 3σ detection of the cross-correlation signal in the optimistic scenario, while CMB-S4 will be able to detect the signal with S/N ratio $\gtrsim 10$ even in more pessimistic scenarios.

The next generation of CMB surveys will map the spatial fluctuations of free electrons out to high redshifts, informing models of reionization and providing clues on the gas physics around galaxies [57,58]. Moving forward, a compelling avenue to explore would be to extend our modeling to provide direct constraints on phenomenological or physical reionization parameters, or the joint analysis with observables probing other phases of matter, such as Compton-y or 21-cm maps (e.g., [21,59,60]), or the galaxy distribution itself.

APPENDIX: A HALO MODEL FOR REIONIZATION

In this appendix, we briefly review the basics of the analytic halo model formalism we use to predict the power spectrum of the optical depth field ($\tau\tau$) and its cross-correlation with CMB lensing potential ($\phi\tau$).

The model presented here is largely based on the phenomenological parametrization introduced in [14,18,40]. In this picture, dark matter halos host the HII regions surrounding the ionizing sources such as galaxies or quasars. These regions are modeled as fully ionized spherical bubbles of radius R inside the neutral intergalactic medium. As time progresses, the ionized bubbles grow and eventually overlap with each other until the Universe is fully reionized.

To calculate the optical depth autospectrum $C_\ell^{\tau\tau}$ we need a model for the three-dimensional power spectrum of the density-weighted ionization fraction fluctuations, $P_{\delta_e \delta_e}(k, z)$. This quantity captures the physics and morphology of reionization which sensitively depend on the size distribution and evolution of the ionized bubbles, which is the first ingredient of the model.

Assuming that the bubble sizes are drawn from a log-normal distribution with characteristic size $\bar{R}$ and variance given by $\sigma_{\ln R}^2$ (e.g., [61–63]), we can write

$$P(R) = \frac{1}{R} \frac{1}{\sqrt{2\pi\sigma_{\ln R}^2}} e^{-[\ln(R/\bar{R})]^2/(2\sigma_{\ln R}^2)}, \qquad (A1)$$

so that the average bubble volume can be calculated as

$$\langle V_b \rangle \equiv \int dR\, V_b(R) P(R) = \frac{4\pi\bar{R}^3}{3} e^{9\sigma_{\ln R}^2/2}. \quad (A2)$$

While, in principle, both $\bar{R}$ and $\sigma_{\ln R}^2$ are functions of z, and would require complex radiative transfer simulations to be evaluated, we follow the previous works and assume that they do not evolve over cosmic time.

The model further assumes that the probability that a given point in space $\mathbf{r}$ is ionized is determined by a Poisson process with fluctuating mean,

$$\langle x_e(\mathbf{r}) \rangle_P = 1 - e^{-n_b(\mathbf{r})\langle V_b \rangle}, \qquad (A3)$$

where $n_b(\mathbf{r})$ is the density of HII bubbles, and the brackets $\langle \ldots \rangle_P$ denote averaging over the Poisson process.

The HII bubbles are then assumed to trace the large-scale structure with some bias b,

$$n_b(\mathbf{r}) = \bar{n}_b(1 + b\delta_W(\mathbf{r})), \qquad (A4)$$

where $\delta_W = \int d^3\mathbf{r}' \delta(\mathbf{r}) W_R(\mathbf{r} - \mathbf{r}')$ is the matter overdensity δ smoothed by a top-hat window of radius R which in Fourier space is given by

$$W_R(k) = \frac{3}{(kR)^3}[\sin(kR) - kR\cos(kR)]. \qquad (A5)$$

For simplicity, we assume that the bubble bias does not evolve with redshift and is independent of the bubble size. Following previous works (e.g., [18,56]), in this paper we fix $b = 6$. With these definitions, the average bubble

number density can be related to the mean ionization fraction as $\bar{n}_b = -\ln(1 - \bar{x}_e)/\langle V_b \rangle$. For completeness, we further define the two-bubble and one-bubble window functions averaged over the bubble radius distributions respectively as follows:

$$\langle W_R \rangle(k) = \frac{1}{\langle V_b \rangle} \int_0^\infty dR \, P(R) V_b(R) W_R(kR), \quad (A6)$$

$$\langle W_R^2 \rangle(k) = \frac{1}{\langle V_b \rangle^2} \int_0^\infty dR \, P(R) [V_b(R) W_R(kR)]^2. \quad (A7)$$

In analogy with the standard halo model of large-scale structure (e.g., [64]), we can decompose the HII power spectrum into the sum of a 1-bubble term, which dominates at scales $r \ll R$, and a 2-bubble contribution that prevails at larger scales, $r \gg R$: $P_{\delta_e \delta_e}(k) = P_{\delta_e \delta_e}^{1b}(k) + P_{\delta_e \delta_e}^{2b}(k)$.[3] The analytic expressions of the 2-bubble and 1-bubble terms of the ionized hydrogen power spectrum are given, respectively, by

---

[3]Here we suppress the redshift dependence for brevity.

$$P_{\delta_e \delta_e}^{2b}(k) = [(1 - \bar{x}_e) \ln(1 - \bar{x}_e) b \langle W_R \rangle(k) - \bar{x}_e]^2 P_{\delta\delta}(k), \quad (A8)$$

$$P_{\delta_e \delta_e}^{1b}(k) = \bar{x}_e (1 - \bar{x}_e)[\langle V_b \rangle \langle W_R^2 \rangle(k) + \tilde{P}(k)], \quad (A9)$$

where $\tilde{P}(k)$ can be approximated as [14,40]

$$\tilde{P}(k) \approx \frac{P_{\delta\delta}(k) \langle V_b \rangle \langle \sigma_R^2 \rangle}{\sqrt{P_{\delta\delta}^2(k) + \langle V_b \rangle^2 \langle \sigma_R^2 \rangle^2}}, \quad (A10)$$

and $\langle \sigma_R^2 \rangle$ is the matter density variance top hat smoothed over the bubble radius distribution. In the formulae above, $P_{\delta\delta}(k)$ denotes the linear matter power spectrum.

Similarly, the 2- and 1-bubble terms relative to the cross-power spectrum between the ionization fraction fluctuations and the matter density contrast can be evaluated as

$$P_{\delta\delta_e}^{2b} = [\bar{x}_e - (1 - \bar{x}_e) \ln(1 - \bar{x}_e) b \langle W_R \rangle(k)] P_{\delta\delta}(k) \quad (A11)$$

$$P_{\delta\delta_e}^{1b} \approx 0, \quad (A12)$$

where the term $P_{\delta\delta_e}^{1b}$ is zero since we assume the bubbles are completely ionized.

[1] N. Aghanim, S. Majumdar, and J. Silk, Rep. Prog. Phys. **71,** 066902 (2008).

[2] J. E. Carlstrom, G. P. Holder, and E. D. Reese, Annu. Rev. Astron. Astrophys. **40,** 643 (2002).

[3] A. Lewis and A. Challinor, Phys. Rep. **429,** 1 (2006).

[4] R. A. Sunyaev and Y. B. Zel'dovich, Astrophys. Space Phys. Rev. **2,** 66 (1970).

[5] R. A. Sunyaev and Ya. B. Zeldovich, Astrophys. Space Phys. Rev. **4,** 173 (1972).

[6] N. Aghanim, Y. Akrami et al. (Planck Collaboration), Astron. Astrophys. **641,** A6 (2020).

[7] C. L. Reichardt, S. Patil, P. A. R. Ade et al., Astrophys. J. **908,** 199 (2021).

[8] S. Aiola, E. Calabrese, L. Maurin et al., arXiv:2007.07288.

[9] P. Ade, J. Aguirre et al. (Simons Observatory Collaboration), J. Cosmol. Astropart. Phys. 02 (2019) 056.

[10] M. Aravena, J. Austermann, K. Basu et al., arXiv:1909.02587.

[11] CMB-S4 Collaboration, arXiv:1907.04473.

[12] C. L. Chang et al., arXiv:2203.07638.

[13] A. Gruzinov and W. Hu, Astrophys. J. **508,** 435 (1998).

[14] M. J. Mortonson and W. Hu, Astrophys. J. **657,** 1 (2007).

[15] W. Hu, M. M. Hedman, and M. Zaldarriaga, Phys. Rev. D **67,** 043004 (2003).

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

[17] M. Zaldarriaga and U. Seljak, Phys. Rev. D **59,** 123507 (1999).

[18] C. Dvorkin and K. M. Smith, Phys. Rev. D **79,** 043003 (2009).

[19] V. Gluscevic, M. Kamionkowski, and D. Hanson, Phys. Rev. D **87,** 047303 (2013).

[20] T. Namikawa, Phys. Rev. D **97,** 063505 (2018).

[21] T. Namikawa, A. Roy, B. D. Sherwin, N. Battaglia, and D. N. Spergel, Phys. Rev. D **104,** 063514 (2021).

[22] M. Su, A. P. S. Yadav, M. McQuinn, J. Yoo, and M. Zaldarriaga, arXiv:1106.4313.

[23] S. Adachi et al., Phys. Rev. Lett. **124,** 131301 (2020).

[24] M. Millea et al., Astrophys. J. **922,** 259 (2021).

[25] J. Carron and A. Lewis, Phys. Rev. D **96,** 063510 (2017).

[26] C. M. Hirata and U. Seljak, Phys. Rev. D **67,** 043001 (2003).

[27] C. M. Hirata and U. Seljak, Phys. Rev. D **68,** 083002 (2003).

[28] M. Millea, E. Anderes, and B. D. Wandelt, Phys. Rev. D **100,** 023509 (2019).

[29] M. Millea, E. Anderes, and B. D. Wandelt, Phys. Rev. D **102,** 123542 (2020).

[30] M. Millea and U. Seljak, Phys. Rev. D **105,** 103531 (2022).

[31] J. a. Caldeira, W. L. K. Wu, B. Nord et al., Astron. Comput. **28,** 100307 (2019).

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

[33] C. Feng and G. Holder, Phys. Rev. D **99,** 123502 (2019).

[34] S. Mukherjee, S. Paul, and T. R. Choudhury, Mon. Not. R. Astron. Soc. **486,** 2042 (2019).

[35] A. Roy, G. Kulkarni, P. D. Meerburg, A. Challinor, C. Baccigalupi, A. Lapi, and M. G. Haehnelt, J. Cosmol. Astropart. Phys. 01 (2021) 003.

[36] A. Lewis, Phys. Rev. D **78**, 023002 (2008).

[37] G. Holder, I. T. Iliev, and G. Mellema, Astrophys. J. Lett. **663**, L1 (2007).

[38] D. N. Limber, Astrophys. J. **117**, 134 (1953).

[39] S. Furlanetto, M. Zaldarriaga, and L. Hernquist, Astrophys. J. **613**, 1 (2004).

[40] X. Wang and W. Hu, Astrophys. J. **643**, 585 (2006).

[41] N. Battaglia, A. Natarajan, H. Trac, R. Cen, and A. Loeb, Astrophys. J. **776**, 83 (2013).

[42] O. Dore, G. Holder, M. Alvarez, I T. Iliev, G. Mellema, U -L Pen, and P. R. Shapiro, Phys. Rev. D **76**, 043002 (2007).

[43] C. Dvorkin, W. Hu, and K. M. Smith, Phys. Rev. D **79**, 107302 (2009).

[44] S. Paul, S. Mukherjee, and T. R. Choudhury, Mon. Not. R. Astron. Soc. **500**, 232 (2020).

[45] T. R. Choudhury, S. Mukherjee, and S. Paul, Mon. Not. R. Astron. Soc. **501**, L7 (2021).

[46] A. Gorce, S. Ilić, M. Douspis, D. Aubert, and M. Langer, Astron. Astrophys. **640**, A90 (2020).

[47] M. G. Santos, A. Cooray, Z. Haiman, L. Knox, and C.-P. Ma, Astrophys. J. **598**, 756 (2003).

[48] O. Zahn, M. Zaldarriaga, L. Hernquist, and M. McQuinn, Astrophys. J. **630**, 657 (2005).

[49] A. Natarajan, N. Battaglia, H. Trac, U. L. Pen, and A. Loeb, Astrophys. J. **776**, 82 (2013).

[50] S. Raghunathan, Astrophys. J. **928**, 16 (2022).

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

[52] J. Delabrouille, J. F. Cardoso, M. Le Jeune, G. Fay, and F. Guilloux, Astron. Astrophys. **493**, 835 (2009).

[53] M. Tegmark and G. Efstathiou, Mon. Not. R. Astron. Soc. **281**, 1297 (1996).

[54] M. Millea, arXiv:2209.10512.

[55] U. Seljak, G. Aslanyan, Y. Feng, and C. Modi, J. Cosmol. Astropart. Phys. 12 (2017) 009.

[56] A. Roy, A. Lapi, D. Spergel, and C. Baccigalupi, J. Cosmol. Astropart. Phys. 05 (2018) 014.

[57] K. Abazajian et al., arXiv:2203.08024.

[58] J. Delabrouille et al., Exp. Astron. **51**, 1471 (2021).

[59] P. D. Meerburg, C. Dvorkin, and D. N. Spergel, Astrophys. J. **779**, 124 (2013).

[60] A. Roy, A. van Engelen, V. Gluscevic, and N. Battaglia, arXiv:2201.05076.

[61] M. M. Friedrich, G. Mellema, M. A. Alvarez, P. R. Shapiro, and I. T. Iliev, Mon. Not. R. Astron. Soc. **413**, 1353 (2011).

[62] Y. Lin, S. P. Oh, S. R. Furlanetto, and P. M. Sutter, Mon. Not. R. Astron. Soc. **461**, 3361 (2016).

[63] O. Zahn, A. Lidz, M. Mc Quinn et al., Astrophys. J. **654**, 12 (2007).

[64] A. Cooray and R. Sheth, Phys. Rep. **372**, 1 (2002).