Maximum likelihood kinetic Sunyaev-Zel'dovich velocity reconstruction

Dagoberto Contreras,[1] Fiona McCarthy,[2,3,4] and Matthew C. Johnson[1,3]

[1]*Department of Physics and Astronomy, York University, Toronto, Ontario M3J 1P3, Canada*
[2]*Center for Computational Astrophysics, Flatiron Institute,*
162 5th Avenue, New York, New York 10010 USA
[3]*Perimeter Institute for Theoretical Physics, Waterloo, Ontario N2L 2Y5, Canada*
[4]*Department of Physics and Astronomy, University of Waterloo, Waterloo, Ontario N2L 3G1, Canada*

The kinetic Sunyaev Zel'dovich (kSZ) effect, cosmic microwave background (CMB) temperature anisotropies induced by the scattering of CMB photons from free electrons, will be measured by near-term CMB experiments at high significance. By combining CMB temperature anisotropies with a tracer of structure, such as a galaxy redshift survey, previous literature introduced a number of techniques to reconstruct the radial velocity field. This reconstructed radial velocity field encapsulates the majority of the cosmological information contained in the kSZ temperature anisotropies, and can provide powerful new tests of the standard cosmological model and theories beyond it. In this paper, we introduce a new estimator for the radial velocity field based on a coarse-grained maximum likelihood fit for the kSZ component of the temperature anisotropies, given a tracer of the optical depth. We demonstrate that this maximum likelihood estimator yields a higher fidelity reconstruction than existing quadratic estimators in the low-noise and high-resolution regime targeted by upcoming CMB experiments. We describe implementations of the maximum likelihood estimator in both harmonic space and map space, using either direct measurements of the optical depth field or a galaxy survey as a tracer. We comment briefly on the impact of biases introduced by imperfect reconstruction of the optical depth field.

I. INTRODUCTION

Future cosmic microwave background (CMB) experiments such as Simons Observatory [1], CMB-S4 [2], and CMB-HD [3] will probe the high-resolution, low-noise frontier where secondary CMB anisotropies dominate. Secondary anisotropies arise due to the electromagnetic (Sunyaev Zel'dovich effect) or gravitational (lensing) scattering of CMB photons from structure, and encode a significant amount of astrophysical and cosmological information. How to access this information is an active area of investigation. In this paper, we focus on the kinetic Sunyaev Zel'dovich (kSZ) effect, blackbody temperature anisotropies induced by the scattering of CMB photons from free electrons in bulk motion along the line of sight [4]. We introduce a new estimator for the radial velocity field based on a coarse-grained maximum likelihood solution constructed from the CMB temperature anisotropies and a tracer of the optical depth.

Previous literature introduced a set of quadratic estimators (QEs) for the radial velocity field [5–7] based on the CMB temperature anisotropies and a galaxy survey as a tracer of the optical depth.[1] Fisher forecasts [5–7], the

analysis of N-body simulations [9,10], and the analysis of Gaussian mocks containing various foregrounds and systematics [7] indicate that these QEs will yield high-fidelity reconstructions of the radial velocity field with near-term CMB and galaxy surveys. The QE reconstructs the radial velocity field from statistical anisotropies in the cross-power spectrum of the CMB temperature and a redshift-binned galaxy survey. The estimator relies on the assumption that all the fields involved are Gaussian, and is based on the theoretical ensemble average spectra, which are used to find the appropriate weights for the estimators to be unbiased and have the minimum variance. However, these properties of the QE lead to significant signal-to-noise penalties and biases in the high signal-to-noise regime. For example, the non-Gaussian nature of the optical depth and galaxy density on small scales leads to the so-called $N^{(3/2)}$ bias which can exceed the Gaussian reconstruction noise in the high signal-to-noise regime [10]. For sufficiently low-noise and high-resolution CMB experiments, and in the absence of foregrounds, the dominant contribution to the observed temperature anisotropies on small angular scales will be the kSZ effect itself. In this limit, the QE is limited by sample variance, since it relies on searching for statistical anisotropies beyond those expected in the ensemble average. However, because one is able to accurately measure the detailed kSZ temperature anisotropies in this limit, it should be possible to do better

---

[1]More precisely, one reconstructs the remote dipole field: the CMB dipole seen in the local rest frame of the scatterer, projected along the line of sight. The largest contribution to the remote dipole field is given by the radial peculiar velocity, but there are additional "primordial" contributions; see Refs. [5,8] for discussion.

by using the properties of our realization rather than the ensemble statistics (e.g. to use a form of sample variance cancellation).

To overcome these two shortcomings of the QE, we introduce a new estimator for the radial velocity field based on a coarse-grained maximum likelihood condition. This maximum likelihood (MaxL) estimator does not rely on the assumption that the underlying fields are Gaussian (hence we do not expect the analog of the $N^{(3/2)}$ bias), and relies on the realization of the underlying fields rather than their statistical properties (and hence should incorporate a form of sample variance cancellation). We therefore expect the MaxL estimator to have advantages over the QE in the high signal-to-noise regime, where the kSZ effect is signal dominated (e.g. over the primary CMB, instrumental noise, and foregrounds) in the CMB on small angular scales. Because it relies on the properties of a realization, forecasting with the MaxL estimator requires simulations, unlike the QE, whose statistical properties can be calculated analytically. We create properly correlated random Gaussian realizations of the underlying fields and make direct comparisons between the reconstruction of the QE as implemented previously [7] and the MaxL estimator. Because the realizations are purely Gaussian, we are unable to demonstrate the absence of a non-Gaussian bias for the MaxL estimator, and therefore it will be important in future work to compare the results of applying the MaxL estimator and QE to N-body simulations.

A key ingredient of the MaxL estimator is an estimate of the optical depth field. In the absence of a direct measurement (e.g. by using a large number of fast radio bursts [11]), we must infer the optical depth from a tracer of large scale structure, such as a galaxy survey. This inference will be imperfect, leading to a multiplicative optical depth bias on the reconstructed velocities [6,11–13]. Further multiplicative biases will be introduced by redshift errors in the galaxy survey and various systematics in the galaxy and CMB surveys. These multiplicative biases on the QE are described in detail in Ref. [7], and in this paper we describe their nature using a simplistic estimator for the optical depth field. Our implementation of the MaxL estimator can likely be improved upon through the use of more accurate methods to estimate the optical depth field, e.g. using machine learning techniques trained on accurate hydrodynamical simulations (see e.g. [14]).

The plan of the paper is as follows. In Sec. II we provide some intuition and motivation for the MaxL estimator. In Sec. III we review our model for the kSZ component of the temperature anisotropies, which is largely a review of Ref. [7]. In Sec. IV, we derive an implementation of the maximum likelihood estimator (given the optical depth field) for the radial velocity field in both harmonic space and map space. In Sec. V we compare the performance of the MaxL estimator and QE on a set of correlated Gaussian mock datasets. In Sec. VI we discuss inference of the

optical depth field from a galaxy redshift survey and the associated optical depth bias on the reconstructed velocity; we conclude in Sec. VII.

A. Some notation

Throughout the paper, we refer to objects defined on a sphere in both "harmonic" and "map" space. We label the multipole moments of fields in harmonic space by the pairs $(\ell, m), (\ell', m')\dots$ and of the pixels in map space by the letters $i, j, \dots$. For coarse-grained quantities, we use capital letters $(L, M), (L', M'), \dots$ or $I, J\dots$. We denote operators in boldface, such that τ is a field but $\boldsymbol{\tau}$ is an operator.

II. INTUITION AND MOTIVATION

In this section we give a heuristic description of the QE that recovers the radial velocity field from a temperature and optical depth measurement, and motivate a new estimator that can avoid some of the QE reconstruction noise in the high-signal-to-noise regime of velocity reconstruction.

The kSZ contribution to the CMB temperature Θ is of the form τv, where v is the radial velocity field and τ is the optical depth field, such that

$$\Theta = \Theta^{\mathrm{pCMB}} + \Theta^{\mathrm{kSZ}},$$
$$\sim \Theta^{\mathrm{pCMB}} + \tau v, \tag{1}$$

where Θ^{pCMB} describes all contributions to the temperature anisotropies that are not kSZ (this includes primary CMB contributions as well as other foreground contributions and instrumental noise). Here, we employ a toy model where kSZ is sourced by a two-dimensional field; later we include line-of-sight effects. The goal is to use a measurement of Θ and τ to recover v. The QE recovers the velocity from this combination and their statistics:

$$\hat{v}^{\mathrm{QE}} \sim \frac{1}{A}\Theta\tau, \tag{2}$$

where A is a theoretically calculated filter that can be thought of as $\sim\langle\tau^2\rangle$; it is calculated from the ensemble statistics of τ. The QE differs from the true velocity by

$$\hat{v}^{\mathrm{QE}} - v \sim \frac{1}{\langle\tau^2\rangle}\Theta\tau - v \sim v\left(\frac{\tau^2}{\langle\tau^2\rangle} - 1\right) + \frac{\Theta^{\mathrm{pCMB}}}{\langle\tau^2\rangle}. \tag{3}$$

In particular, there is a contribution to the residual sourced by the difference of the realization of τ^2 from its theoretically expected variance. This contribution persists even when Θ^{pCMB} is vanishingly small.

However, in the limit where the Θ^{kSZ} contributions dominate over the non-kSZ contributions, we can avoid this contribution to the reconstruction noise by neglecting Θ^{pCMB} and simply "inverting" Eq. (1):

$$\hat{v} \sim \frac{\Theta}{\tau}. \tag{4}$$

In this case, the "reconstruction noise" only has contributions from the non-kSZ contributions to the temperature:

$$\hat{v} - v \sim \frac{\Theta^{\text{pCMB}}}{\tau}. \tag{5}$$

This allows us to make a lower-noise estimate of v by utilizing angular scales where the kSZ temperature dominates over the other contributions to the CMB temperature, similar to what is done in the gradient-inversion approach to CMB lensing reconstruction [15].

In the remainder of this paper, we develop this idea more rigorously by deriving a new estimator from the maximization of a likelihood, and compare explicitly the performance of the "MaxL" estimator with that of the QE. Note that further complication arises from the fact that the kSZ temperature is not *truly* sourced by the product of one velocity and one electron density field, as in Eq. (1), but by the integral of this product along the line of sight. We will deal with this explicitly by "redshift-binning" the velocity and optical depth fields, with the end goal being to reconstruct the redshift-binned velocity field.

III. THE KINETIC SUNYAEV ZEL'DOVICH EFFECT

The observed kSZ temperature anisotropies in the direction $\hat{n}$ are given by the line-of-sight integral over radial comoving distance χ

$$\Theta^{\text{kSZ}}(\hat{n}) = \int_0^{\chi_{\text{max}}} d\chi \, \dot{\tau}(\hat{n}, \chi) v(\hat{n}, \chi). \tag{6}$$

Here, $v(\hat{n}, \chi)$ is the radial velocity field and the differential optical depth $\dot{\tau}(\hat{n}, \chi)$ is

$$\dot{\tau}(\hat{n}, \chi) = -\sigma_T a(\chi) \bar{n}_e(\chi)(1 + \delta_e(\hat{n}, \chi)), \tag{7}$$

with σ_T the Thompson cross section, $a(\chi)$ the scale factor, $\bar{n}_e(\chi)$ the cosmological mean electron density, and $\delta_e(\hat{n}, \chi)$ the electron overdensity field. There are contributions to the kSZ signal from during and after the epoch of reionization. We are primarily concerned with the late-time kSZ effect in this paper,[2] and so in Eq. (6) we introduce a cutoff χ_{max} on the line-of-sight integral at some time after reionization (which also justifies our neglect of the optical depth factor $e^{-\tau}$ in this equation).

Following Ref. [7], we decompose fields along the radial direction in the basis of Haar wavelets:

---

[2]See Ref. [16] for a discussion of kSZ velocity reconstruction during the epoch of reionization.

$$v(\hat{n}, \chi) = \sum_{s=0}^{\infty} v^s(\hat{n}) h^s(\chi), \tag{8}$$

where h^s are the Haar wavelet basis functions; these form a complete basis on the interval $0 \leq \chi \leq \chi_{\text{max}}$. The Haar expansion has the additional useful property that if we truncate at a fixed index N, the sum in the Haar basis is simply related to a sum over radial top-hat bin averages over bins of equal comoving size $\Delta\chi = \chi_{\text{max}}/N$:

$$v(\hat{n}, \chi) = \sum_{\alpha=0}^{N-1} v^\alpha(\hat{n})\Pi^\alpha(\chi) + \sum_{s=N}^{\infty} v^s(\hat{n})h^s(\chi), \tag{9}$$

where Π^α are radial top-hat bins. We will reserve Greek superscripts $\alpha, \beta, \gamma, \ldots$ to denote the top-hat bin averages. The bin averages v^α give a coarse-grained description of the field on the past light cone, and the remaining sum over Haar moments gives the residual fine-grained information. We can therefore write the line-of-sight integral in Eq. (6) as the sum:

$$\begin{aligned} \Theta^{\text{kSZ}}(\hat{n}) &= \sum_{s=0}^{\infty} \tau^s(\hat{n}) v^s(\hat{n}) \\ &= \sum_{\alpha=0}^{N-1} \tau^\alpha(\hat{n}) v^\alpha(\hat{n}) + \sum_{s=N}^{\infty} \tau^s(\hat{n}) v^s(\hat{n}), \end{aligned} \tag{10}$$

where we have defined $\tau^\alpha(\hat{n}) \equiv \Delta\chi \dot{\tau}^\alpha(\hat{n})$. Below, we approximate the continuum limit by truncating at fixed N. For the contribution to kSZ from the redshift range $0 < z < 5$ most of the kSZ autopower is captured for $N = 512$ as shown in Ref. [7]. However, in practice we will have to truncate the sum at a somewhat lower value of N for the application of the estimators derived below.

It will be useful to think of the optical depth as a set of N operators that map the bin-averaged radial velocity field to the observed temperature anisotropies, in which case we can write the sky-basis-independent equation truncated at fixed N:

$$\Theta^{\text{kSZ}} \simeq \sum_{\alpha=0}^{N-1} \boldsymbol{\tau}^\alpha \cdot v^\alpha. \tag{11}$$

Here, just as any field defined on a sphere, Θ^{kSZ} can be represented either in map space or harmonic space. In map space, Θ_i has dimension equal to the number of pixels N_{pix} at some resolution (e.g. using the Healpix pixelization scheme [17]); v_i^α is a set of N fields each with dimension equal to N_{pix}; τ_{ij}^α is a set of N operators each with dimension equal to $N_{\text{pix}} \times N_{\text{pix}}$. In harmonic space, $\Theta_{\ell m}$ has dimension equal to the number of modes required to describe the field at the appropriate resolution (which should in principle be equal to N_{pix}, but in practice is

somewhat larger); we denote this number as $N_{\rm harm}$. In Appendix A, we present the form of $\boldsymbol{\tau}$ in each representation, as well as a mixed harmonic-map space representation.

IV. MAXIMUM LIKELIHOOD ESTIMATOR

The dominant blackbody components of the observed CMB temperature anisotropies are the lensed primary CMB and the late-time and reionization kSZ. Neglecting the reionization kSZ, the blackbody component of the temperature anisotropy is given by

$$\Theta = \Theta^{\rm pCMB} + \sum_{s=0}^{\infty} \tau^s \cdot v^s. \qquad (12)$$

Here $\Theta^{\rm pCMB}$ are the primary CMB anisotropies in map or harmonic space, which we assume purely Gaussian with covariance $\boldsymbol{C}^{\rm pCMB}$. The likelihood function $\mathcal{L}$ for the primary CMB is then

$$\ln \mathcal{L} = \left(\Theta - \sum_{s=0}^{\infty} \boldsymbol{\tau}^s \cdot v^s\right)^{\dagger} (\boldsymbol{C}^{\rm pCMB})^{-1} \left(\Theta - \sum_{s=0}^{\infty} \boldsymbol{\tau}^s \cdot v^s\right)$$
$$+ \ln \, \det \boldsymbol{C}^{\rm pCMB}. \qquad (13)$$

Extremizing $\mathcal{L}$ with respect to $(v^q)^{\dagger}$:

$$\frac{\delta}{\delta(v^q)^{\dagger}} \ln \mathcal{L} = (\boldsymbol{\tau}^q)^{\dagger}(\boldsymbol{C}^{\rm pCMB})^{-1}\left(\Theta - \sum_{s=0}^{\infty} \boldsymbol{\tau}^s \cdot v^s\right) = 0$$

$$\Rightarrow (\boldsymbol{\tau}^q)^{\dagger}(\boldsymbol{C}^{\rm pCMB})^{-1}\Theta = \sum_{s=0}^{\infty}[(\boldsymbol{\tau}^q)^{\dagger}(\boldsymbol{C}^{\rm pCMB})^{-1}\boldsymbol{\tau}^s] \cdot v^s$$

$$(\boldsymbol{\tau}^{\alpha})^{\dagger}(\boldsymbol{C}^{\rm pCMB})^{-1}\Theta \simeq \sum_{\beta=0}^{N-1}[(\boldsymbol{\tau}^{\alpha})^{\dagger}(\boldsymbol{C}^{\rm pCMB})^{-1}\boldsymbol{\tau}^{\beta}] \cdot v^{\beta}. \qquad (14)$$

In the last line we truncate the radial sum at N and work with the bin coefficients (denoted by the Greek indices). Note that Θ and $\boldsymbol{\tau}$ are fixed, and our goal is to find the set of v^{α} that maximizes the likelihood. Organizing both v^{β} and $(\boldsymbol{\tau}^{\alpha})^{\dagger}(\boldsymbol{C}^{\rm pCMB})^{-1}\Theta$ into $(N_{\rm pix} \times N)$-dimensional vectors containing both the bin and pixel information, and organizing $[(\boldsymbol{\tau}^{\alpha})^{\dagger}(\boldsymbol{C}^{\rm pCMB})^{-1}\boldsymbol{\tau}^{\beta}]$ into an $(N_{\rm pix}N) \times (N_{\rm pix}N)$-dimensional operator, Eq. (14) is a linear equation for the velocity field:

$$[\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\rm pCMB})^{-1}\Theta] = [\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\rm pCMB})^{-1}\boldsymbol{\tau}] \cdot v. \qquad (15)$$

This equation will have a solution when the operator $\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\rm pCMB})^{-1}\boldsymbol{\tau}$ has a well-defined inverse. In general, this is not the case, since there will be many highly oscillatory functions v that nearly satisfy the maximum likelihood condition and pixels at which the operator (whose eigen-basis is in map space) is nearly zero (indicating small eigenvalues). However, we can first perform the matrix

products in $\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\rm pCMB})^{-1}\boldsymbol{\tau}$ at full resolution, and then coarse-grain this operator so that it is invertible. The maximum-likelihood condition is then

$$W([\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\rm pCMB})^{-1}\Theta]) = W([\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\rm pCMB})^{-1}\boldsymbol{\tau}]) \cdot W(v)$$
$$\Rightarrow W(v) = [W([\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\rm pCMB})^{-1}\boldsymbol{\tau}])]^{-1}$$
$$\times W([\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\rm pCMB})^{-1}\Theta]), \qquad (16)$$

where $W(\cdot)$ is the coarse-graining procedure; we demonstrate explicitly below that it is appropriate to separately coarse-grain $[\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\rm pCMB})^{-1}\boldsymbol{\tau}]$ and v. The proposal is that with Eq. (16), we can reconstruct the coarse-grained velocity $W(v)$. We now separately examine if there is a useful solution to the maximum likelihood condition in harmonic or map space for the velocity field.

A. Harmonic space estimator

Before coarse-graining, the maximum likelihood condition in harmonic space is

$$[\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\rm pCMB})^{-1}\Theta]_{\ell m}^{\alpha} = \sum_{\ell'm';\beta}[\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\rm pCMB})^{-1}\boldsymbol{\tau}]_{\ell m;\ell'm'}^{\alpha\beta} v_{\ell'm'}^{\beta}. \qquad (17)$$

In Appendix A we evaluate the components of the operators and vectors in Eq. (17) explicitly, obtaining

$$[\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\rm pCMB})^{-1}\boldsymbol{\tau}]_{\ell m;\ell'm'}^{\alpha\beta} = \int d\Omega \tau^{\alpha}(\hat{n})\bar{\tau}^{\beta}(\hat{n})Y_{\ell m}^{*}(\hat{n})Y_{\ell'm'}(\hat{n}),$$
$$(18)$$

and

$$[\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\rm pCMB})^{-1}\Theta]_{\ell m}^{\beta} = \int d\Omega \tau^{\beta}(\hat{n})\bar{\Theta}(\hat{n})Y_{\ell m}^{*}(\hat{n}), \qquad (19)$$

where the bar denotes inverse-variance filtering using $C_{\ell}^{\rm pCMB}$:

$$\bar{A}(\hat{n}) \equiv \sum_{\ell m}\frac{A_{\ell m}}{C_{\ell}^{\rm pCMB}}Y_{\ell m}(\hat{n}). \qquad (20)$$

The coarse-graining operation required to make the operator $[\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\rm pCMB})^{-1}\boldsymbol{\tau}]_{\ell m;\ell'm'}^{\alpha\beta}$ invertible is to cut off the sum at a maximum multipole $L_{\rm max}$, which will be the maximum multipole of the v field that we recover; for more details, see Appendix B. After this coarse-graining, the maximum-likelihood solution for the velocity is

$$\hat{v}_{LM}^{\alpha} = \sum_{L'M'}^{L_{\rm max}}\sum_{\beta=0}^{N-1}([\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\rm pCMB})^{-1}\boldsymbol{\tau}]^{-1})_{LM;L'M'}^{\alpha\beta}$$
$$\times (\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\rm pCMB})^{-1}\Theta)_{L'M'}^{\beta}, \qquad (21)$$

where the $L_{\max}$ cutoff should be done after the calculation of $\boldsymbol{\tau}^\dagger (\boldsymbol{C}^{\mathrm{pCMB}})^{-1}\boldsymbol{\tau}$ but before its inversion. The maximum values N and $L_{\max}$ must be empirically determined from the numerical inversion of $[\boldsymbol{\tau}^\dagger (\boldsymbol{C}^{\mathrm{pCMB}})^{-1}\boldsymbol{\tau}]^{\alpha\beta}_{\ell m;\ell' m'}$ for each realization. A crude approach to finding an approximate solution to the maximum likelihood condition Eq. (17) is to substitute with the ensemble-average operator over realizations of the optical depth, which is explicitly invertible. This approach is presented in Appendix C, where we demonstrate that it yields an expression equivalent to the QE in the low signal-to-noise regime. We note that Eq. (21) is nearly identical to an estimator that was successfully used to reconstruct our peculiar velocity with respect the CMB rest frame using the thermal Sunyaev Zel'dovich effect [18].

We now compute the mean and variance of the MaxL estimator. The estimator Eq. (21) is based on a *fixed realization* of the optical depth, so the statistics of the estimator should be computed over an ensemble of realizations of the primary CMB which we denote by $\langle\rangle_\Theta$. In this ensemble, the mean of the estimator is

$$\langle \hat{v}^\alpha_{LM}\rangle_\Theta = \sum_{L'M'}^{L'_{\max}} \sum_{\beta=0}^{N-1} ([\boldsymbol{\tau}^\dagger (\boldsymbol{C}^{\mathrm{pCMB}})^{-1}\boldsymbol{\tau}]^{-1}]^{\alpha\beta}_{LM;L'M'}$$
$$\times \langle [\boldsymbol{\tau}^\dagger (\boldsymbol{C}^{\mathrm{pCMB}})^{-1}\Theta]^\beta_{L'M'}\rangle_\Theta. \quad (22)$$

We find that

$$\langle \hat{v}^\alpha_{LM}\rangle_\Theta = v^\alpha_{LM} + \beta^\alpha_{LM}, \quad (23)$$

where β is a "coarse-graining" bias

$$\beta^\alpha_{LM} \equiv \sum_{L'M'}^{L_{\max}} \sum_{\gamma=0}^{N-1} \sum_{c=N}^{\infty} \sum_{L''M''}^{L_{\max}} ([\boldsymbol{\tau}^\dagger (\boldsymbol{C}^{\mathrm{pCMB}})^{-1}\boldsymbol{\tau}]^{-1}]^{\alpha\gamma}_{LM;L'M'}$$
$$\times [\boldsymbol{\tau}^\dagger (\boldsymbol{C}^{\mathrm{pCMB}})^{-1}\boldsymbol{\tau}]^{\gamma c}_{L'M';L''M''} v^c_{L''M''}. \quad (24)$$

Details of the calculation that leads to (23) are in Appendix D 1. We also calculate the variance of the estimator, and find that

$$\langle \hat{v}^\alpha_{LM} \hat{v}^\beta_{L'M'}\rangle_\Theta = (C^{vv})^{\alpha\beta}_{LM;L'M'} + N^{\alpha\beta}_{LM;L'M'}, \quad (25)$$

with

$$(C^{vv})^{\alpha\beta}_{LM;L'M'} = (v^\alpha_{LM} + \beta^\alpha_{LM})(v^\beta_{L'M'} + \beta^\beta_{L'M'}), \quad (26)$$

and

$$N^{\alpha\beta}_{LM;L'M'} = ([\boldsymbol{\tau}^\dagger (\boldsymbol{C}^{\mathrm{pCMB}})^{-1}\boldsymbol{\tau}]^{-1})^{\alpha\beta}_{LM;L'M'}. \quad (27)$$

Unlike the QE [5,7], where one must assume that the underlying fields are Gaussian, it is not necessary to make any assumptions about the statistics of the optical depth or velocity fields. Therefore, we do not expect additional contributions to the estimator variance beyond what is presented here. This is in contrast to the QE, where extra contributions to the estimator variance arise from the non-Gaussian nature of the optical depth and velocities [10]. We view this as a significant advantage of the MaxL estimator, since these non-Gaussian contributions are difficult to compute and remove.

Evaluating the harmonic space MaxL estimator in Eq. (21) requires costly, and potentially numerically inaccurate, forward and reverse spherical harmonic transforms. Motivated by this, in the next section we construct a map space estimator. Ultimately, the implementation of the map space estimator is simpler and more accurate, and will be the method used for the forecasts presented below. However, there may be contexts in which the harmonic space estimator is superior, as we comment on below.

B. Map space estimator

We now derive an alternative coarse-graining procedure and obtain a MaxL estimator defined in terms of fields in map space. In map space (as demonstrated in Appendix A), the maximum likelihood condition can be written explicitly as

$$\tau^\alpha(\hat{n})\bar{\Theta}(\hat{n}) = \sum_{\beta=0}^{N-1} \tau^\alpha(\hat{n})\bar{\tau}^\beta(\hat{n})v^\beta(\hat{n}), \quad (28)$$

where the bar denotes inverse-variance filtering by C^{pCMB}_ℓ:

$$\bar{A}(\hat{n}) \equiv \sum_{\ell m} \frac{A_{\ell m}}{C^{\mathrm{pCMB}}_\ell} Y_{\ell m}(\hat{n}). \quad (29)$$

To coarse-grain, we define a set of smoothing kernels $W_I(\hat{n})$ with compact support in regions of the map labeled by $I = 1, 2, \dots I_{\max}$ such that

$$\int d\Omega \, W_I(\hat{n})W_J(\hat{n}) = \frac{\delta_{IJ}}{\Delta\theta^2}, \quad (30)$$

where we have assumed each kernel is nonzero over the same solid angle $\Delta\theta^2$ for all I. A simple choice for $W_I(\hat{n})$ satisfying this criterion, utilized below, is a simple average over equal-area pixels. Note that this is only one choice of definition for $W_I(\hat{n})$ among many possibilities. Multiplying both sides of Eq. (28) by $W_I(\hat{n})$ and integrating over the map, we obtain

$$\int d\Omega \, W_I(\hat{n})\tau^\alpha(\hat{n})\bar{\Theta}(\hat{n}) = \sum_{\beta=0}^{N-1} \int d\Omega \, W_I(\hat{n})\tau^\alpha(\hat{n})\bar{\tau}^\beta(\hat{n})v^\beta(\hat{n}).$$

(31)

Expanding the velocity in terms of an average velocity v_I^β over the region of the map labeled by I and a residual $\Delta v_I^\beta(\hat{n})$ in the same region, we have

$$\int d\Omega W_I(\hat{n})\tau^\alpha(\hat{n})\bar{\Theta}(\hat{n})$$

$$= \sum_{\beta=0}^{N-1} v_I^\beta \int d\Omega W_I(\hat{n})\tau^\alpha(\hat{n})\bar{\tau}^\beta(\hat{n})$$

$$+ \sum_{\beta=0}^{N-1} \int d\Omega W_I(\hat{n})\tau^\alpha(\hat{n})\bar{\tau}^\beta(\hat{n})\Delta v_I^\beta(\hat{n}). \quad (32)$$

Defining the operators

$$\mathcal{W}_{I;\alpha\beta} \equiv \int d\Omega\, W_I(\hat{n})\tau^\alpha(\hat{n})\bar{\tau}^\beta(\hat{n}), \quad (33)$$

we have the following estimator for the average velocity in region I:

$$\hat{v}_I^\alpha = \sum_{\beta=0}^{N-1} \mathcal{W}_{I;\alpha\beta}^{-1} \int d\Omega\, W_I(\hat{n})\tau^\beta(\hat{n})\bar{\Theta}(\hat{n}). \quad (34)$$

In Appendix C, we demonstrate that the QE in the low signal-to-noise regime is recovered upon replacing $\mathcal{W}_{I;\alpha\beta} \to \langle \mathcal{W}_{I;\alpha\beta} \rangle_\tau$.

In our implementation of the map space estimator, we employ the Healpix equal-area pixelization scheme [17]. If we define the input fields at a native Healpix resolution with pixels indexed by $i = 1, 2, \ldots N_{\text{pix}}^{\text{in}}$ and the regions I are defined by pixels at a degraded Healpix resolution $N_{\text{pix}}^{\text{out}}$, then we can write a discrete version of the MaxL estimator:

$$\hat{v}_I^\alpha = \sum_{\beta=0}^{N-1} \left[\sum_{i=1}^{N_{\text{pix}}^{\text{in}}} W_{Ii} \tau_i^\alpha \bar{\tau}_i^\beta \right]^{-1} \left[\sum_{i=1}^{N_{\text{pix}}^{\text{in}}} W_{Ii} \tau_i^\beta \bar{\Theta}_i \right]. \quad (35)$$

Note that W_{Ii} as defined here is simply the Healpix ud_grade function—an operator that averages pixels at a fine resolution labeled by i that fit within a coarse pixel labeled by I (see Ref. [17]).

Taking an ensemble average over realizations of the primary CMB, the estimator mean is

$$\langle \hat{v}_I^\alpha \rangle_\Theta = v_I^\alpha + \beta_I^\alpha \quad (36)$$

where the additive bias β_I^α arises from coarse-graining:

$$\beta_I^\alpha = \sum_{\beta;\gamma=0}^{N-1} \mathcal{W}_{I;\alpha\beta}^{-1} \int d\Omega\, W_I(\hat{n})\tau^\beta(\hat{n})\bar{\tau}^\gamma(\hat{n})\Delta v_I^\gamma(\hat{n})$$

$$+ \sum_{\beta=0}^{N-1} \sum_{c=N}^{\infty} \mathcal{W}_{I;\alpha\beta}^{-1} \int d\Omega\, W_I(\hat{n})\tau^\beta(\hat{n})\bar{\tau}^c(\hat{n})v^c(\hat{n}). \quad (37)$$

The first term is the bias from the pixel-averaging operation and the second term is the bias from coarse-graining along the line of sight. Moving to the estimator variance, we get

$$\langle \hat{v}_I^\alpha \hat{v}_J^\beta \rangle_\Theta = (v_I^\alpha + \beta_I^\alpha)(v_J^\beta + \beta_J^\beta) + \frac{1}{\Delta\theta^2} \mathcal{W}_{I;\alpha\beta}^{-1} \delta_{IJ}. \quad (38)$$

Just as for the harmonic space estimator, the estimator variance includes an additive bias due to coarse-graining as well as a reconstruction noise. Note that the reconstruction noise is local, and does not have any pixel-pixel covariance (although there is bin-bin covariance). A derivation of the estimator mean and variance can be found in Appendix D 2.

The map space estimator is local, only depending on the fields in the region supported by $W_I(\hat{n})$. This allows us to construct I_{max} independent $N \times N$ operators $\mathcal{W}_{I;\alpha\beta}$, which are more easily (and accurately) inverted than the single $(N_{\text{harm}}^2 N) \times (N_{\text{harm}}^2 N)$ operator needed for the harmonic space MaxL estimator. The local nature of the estimator also makes it a convenient choice for masked datasets. Indeed, one could generalize the discussion above to arbitrary complex pixelizations, suited to a particular procedure for masking, foreground removal, or anisotropic noise.

V. IMPLEMENTATION AND FORECASTS

In this section, we describe an implementation of the MaxL estimator and compare its performance to the existing QE introduced in Ref. [5]. Unlike the QE, whose performance can be forecast from theoretical spectra, the performance of the MaxL estimator depends on the realization. We therefore must generate ensembles of simulations to compare the performance of the MaxL estimator to the QE. In this paper, we perform reconstructions on ensembles of properly correlated Gaussian simulations generated using the algorithm of [7]; an analysis of N-body simulations is an important direction for future work. Specifically, we first calculate the theoretical power spectra of the lensed primary CMB $C_\ell^{p\text{CMB}}$; of the velocity field C_ℓ^{vv}; and of the electron density field $C_\ell^{\tau\tau}$. We neglect the cross-power spectrum of electrons and velocity $C_\ell^{\tau v}$ since the estimators utilize the optical depth and velocities on very different scales. We use CAMB [19] to calculate quantities such as the primary CMB and linear matter power spectrum, with the nonlinear regimes modeled with a halo model [7] (see e.g. Ref. [20] for a review of the halo model). The electrons are modeled with the density profiles of [21]. These power spectra are then used to create appropriately correlated Gaussian simulations of the binned velocity v and τ, as well as the primary CMB. To model the non-Gaussian kSZ map, we simply multiply and sum the simulated v and τ maps: $\Theta^{\text{kSZ}} = \sum v^\alpha \tau^\alpha$. We use 32 redshift bins between $z = 0.2$ and $z = 5$, and employ a binning scheme such that each bin has equal comoving width. Further details of this implementation are given in

Sec. III of [7]. We consider the idealized scenario where all non-blackbody foregrounds are absent and where we neglect the non-Gaussianities induced by lensing of the primary CMB. A discussion of these effects in the context of the QE can be found in Ref. [7].

In Fig. 1 we show the ensemble averaged power spectra of the (lensed) primary CMB and of the kSZ effect. Note that we only include power from the 32 redshift bins between $0.2 < z < 5$ used in our analysis, while Ref. [7] demonstrated that approximately 512 radial bins are necessary to capture the majority of the kSZ power in this redshift range. We continue with this model in order to be self-consistent, as using a larger number of bins would be too computationally expensive. We do not expect this approximation to affect our conclusions.

Throughout this section, we consider the case where we have access to a direct measurement of the optical depth. In practice, we will not have access to such a map: one must use a tracer of the electrons such as a galaxy survey. Connecting such a tracer to the optical depth necessitates further modeling, in particular through the cross power spectrum $C_\ell^{\tau g}$. Incorrect modeling of $C_\ell^{\tau g}$ leads to a bias on the reconstruction known as the optical depth bias; we comment on this further in Sec. VI. Additionally, we first consider the case where $C_\ell^{\tau\tau}$ is known to arbitrarily high resolution, and do not consider realistic effects such as the shot noise of a galaxy survey or redshift error on a photometric survey.

We have validated two pipelines, based the map and harmonic space MaxL estimators presented above. We find that the map space estimator has a number of computational advantages. The harmonic space estimator involves several forward and reverse spherical harmonic transforms, which in the Healpix pixelization are not information-preserving. This loss of information leads to numerical inaccuracies in the reconstruction which can be overcome by using exact, but less efficient, pixelization schemes. In addition, the
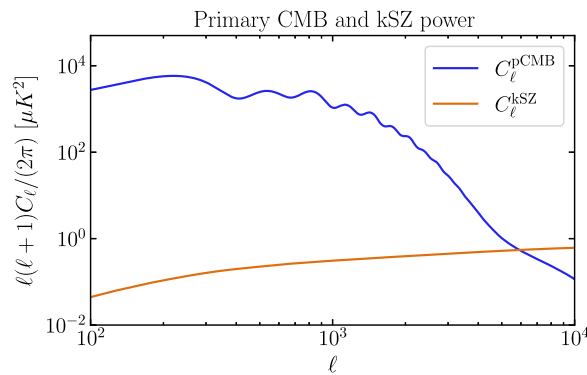
operator inversion required in the velocity estimator is costly when incorporating the full multipole and bin covariance; this places computational limitations on the degree of coarse graining that can be considered. The map space estimator overcomes both of these challenges, and so we focus on this implementation in what follows.

In our first illustration of the MaxL estimator, we completely remove the primary CMB and perform reconstruction on kSZ maps alone (we add a negligible amount of white noise to ensure that we are not dividing by zero when we filter). We use the equal-area Healpix pixelization [17], where the resolution is set by the parameter N_{side}, related to the number of pixels in the map by $N_{\text{pix}} = 12N_{\text{side}}^2$. We apply both the MaxL estimator and the QE to the kSZ-only and optical depth maps for a single realization at a native resolution of $N_{\text{side}}^{\text{in}} = 2048$. Recalling the discussion in Sec. II, we expect the QE to be limited by reconstruction noise, while the MaxL estimator is formally noiseless if the maximum likelihood condition can be inverted without coarse graining. Some degree of coarse graining is necessary for a well-defined solution to the maximum likelihood condition, and we therefore expect the MaxL estimator to be limited by the additive coarse graining bias as described in the previous section.

The power spectrum of the true velocity and the power spectrum of the reconstruction residuals (defined by $\hat{v} - v$) are shown in Fig. 2. For the QE, we perform the reconstruction at a resolution of $N_{\text{side}}^{\text{out}} = 128$; for the MaxL estimator, we perform the reconstruction at a variety of resolutions $N_{\text{side}}^{\text{out}} = 16, 64, 128$. For low-resolution output maps ($N_{\text{side}}^{\text{out}} = 16$) the MaxL estimator in fact performs worse than the QE due to a large coarse-graining bias β, but as we increase $N_{\text{side}}^{\text{out}}$ the MaxL estimator becomes superior as expected. We have explicitly checked that the residual agrees with β as calculated directly from (37).

When we include the primary CMB in the temperature maps, we do not expect the improvement of the MaxL estimator over the QE to be as dramatic. In this case, the MaxL estimator will have both reconstruction noise and an additive coarse-graining bias [see Eq. (38)]. We only expect the MaxL estimator to improve upon the QE when the resolution of the input maps is high enough to probe the regime where the amplitude of the kSZ anisotropies are comparable to those of the primary CMB. To test this expected resolution dependence, we produce input maps at $N_{\text{side}}^{\text{in}} = 1024, 2048$. In Fig. 3 we compare the velocity reconstruction using the quadratic and MaxL estimators for $N_{\text{side}}^{\text{out}} = 16, 64, 128$ using maps with $N_{\text{side}}^{\text{in}} = 2048$. In Fig. 4 we plot the ratio of the total signal to noise for the MaxL estimator and the QE; the signal to noise is defined as

FIG. 1. The primary CMB power spectrum and the kSZ power spectrum using 32 radial bins of equal comoving size between $0.2 < z < 5$. At large-ℓ, the kSZ power dominates over the CMB power. It is only once we access these angular scales that we expect the MaxL estimator to show improvements over the QE.

$$S/N_L = \sqrt{\text{Tr}[C_L^{v^\alpha v^\beta} (N_L^{\hat{v}^\alpha \hat{v}^\beta})^{-1}]}, \qquad (39)$$
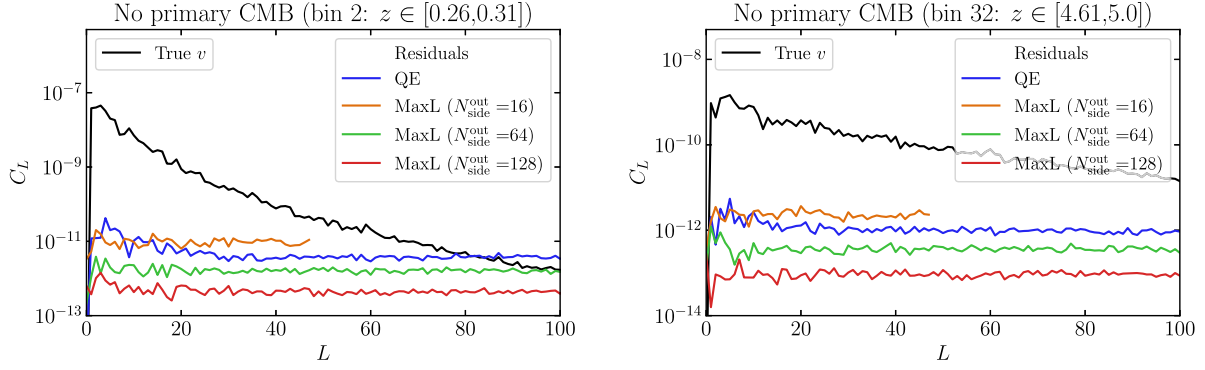
FIG. 2. We show the true velocity power spectrum as well as the power of the residuals $v - \hat{v}^{\mathrm{QE}}$ and $v - \hat{v}^{\mathrm{MaxL}}$. On the left is a low-$z$ bin and on the right is a high-z bin. We show $\hat{v}^{\mathrm{MaxL}}$ for increasing resolution $N_{\mathrm{side}}^{\mathrm{out}}$. At low resolution, the "coarse-graining" bias β is significant and $\hat{v}^{\mathrm{MaxL}}$ performs worse than $\hat{v}^{\mathrm{QE}}$, but as we increase $N_{\mathrm{side}}^{\mathrm{out}}$ this bias decreases and we get a better reconstruction than $\hat{v}^{\mathrm{QE}}$.
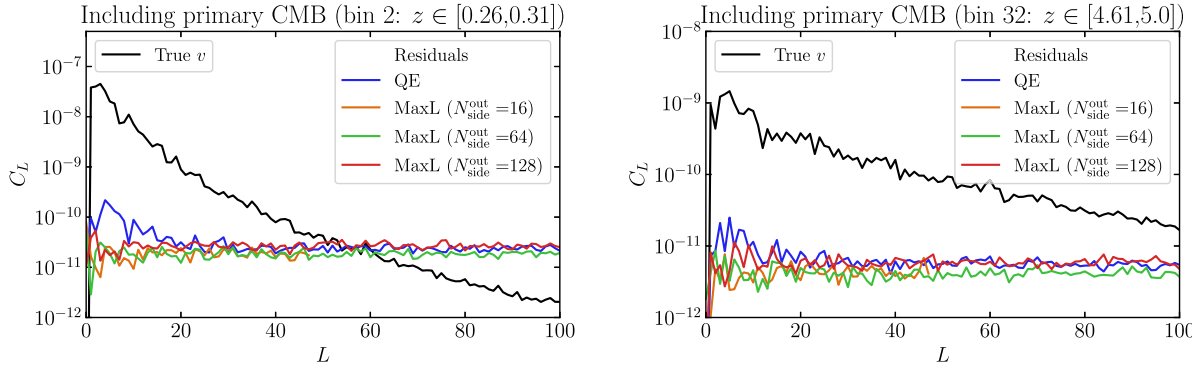

FIG. 3. We show the true velocity power spectrum as well as the power of the residuals $v - \hat{v}^{\mathrm{QE}}$ and $v - \hat{v}^{\mathrm{MaxL}}$. On the left is a low-$z$ bin and on the right is a high-z bin. We show $\hat{v}^{\mathrm{MaxL}}$ for an increasing output resolution. Here we see less of an improvement over the QE than we did in Fig. 2 where we did not include the primary CMB. The MaxL residual power is 75% of the QE residual power on small angular scales, and significantly better on large angular scales.
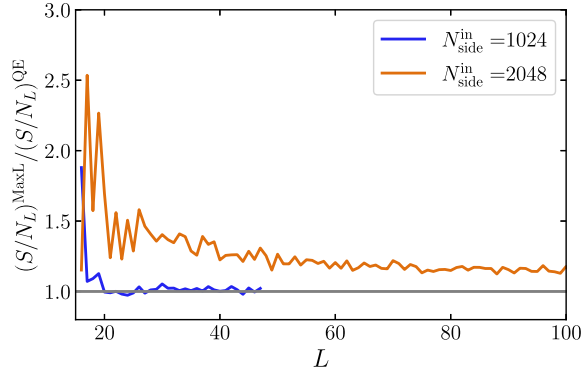

FIG. 4. The ratio of the total signal to noise per mode of the MaxL and quadratic estimators. We find that at $N_{\mathrm{side}}^{\mathrm{in}}$ of 1024, the estimators perform similarly; however, as we go to $N_{\mathrm{side}}^{\mathrm{in}}$ of 2048 and higher, the MaxL estimator has higher signal to noise than the QE. In both cases we have chosen a different $N_{\mathrm{side}}^{\mathrm{out}}$ at which the MaxL performs best: if $N_{\mathrm{side}}^{\mathrm{out}}$ is too low, the coarse-graining bias is too large, and if it is too high, the operator $\mathcal{W}$ becomes uninvertible. These $N_{\mathrm{side}}^{\mathrm{out}}$ corresponded to $N_{\mathrm{side}}^{\mathrm{out}} = \{16, 64\}$ for $N_{\mathrm{side}}^{\mathrm{out}} = \{1024, 2048\}$ respectively.

where $C_L^{v^\alpha v^\beta}$ is the (true) signal matrix and we use the power of the residual as a proxy for the noise $N_L^{\hat{v}^\alpha \hat{v}^\beta}$.

We see that the signal to noise for the MaxL estimator is comparable to that of the QE for $N_{\mathrm{side}}^{\mathrm{in}} = 1024$. Increasing to $N_{\mathrm{side}}^{\mathrm{in}} = 2048$, the MaxL signal to noise per mode is approximately 25% higher than for the QE on small angular scales, and significantly higher on large angular scales. This demonstrates the improvement in performance expected once angular scales on which kSZ power is significant compared to the primary CMB are included.

From these results we see that with perfect knowledge of the optical depth field and including only the primary CMB and kSZ temperature anisotropies, the MaxL estimator does yield a higher fidelity reconstruction than the QE for sufficiently high-resolution data. In the next section, we take one step towards increasing realism and compare the performance of the MaxL and quadratic estimators when the optical depth is inferred from a galaxy survey.

VI. ESTIMATING THE OPTICAL DEPTH

In the absence of a direct measure of the differential optical depth, we must estimate it from a tracer in order to implement the maximum likelihood estimators introduced above. A simple estimator is

$$\hat{\tau}_{\ell m}^{\alpha} = \sum_{\beta,\gamma=0}^{N-1} (C^{\tau g})_{\ell}^{\alpha\beta} ([C^{gg}]^{-1})_{\ell}^{\beta\gamma} g_{\ell m}^{\gamma}, \qquad (40)$$

where $g_{\ell m}^{\gamma}$ is the redshift-binned galaxy density, $(C^{gg})_{\ell}^{\beta\gamma}$ is the galaxy power spectrum, and $(C^{\tau g})_{\ell}^{\alpha\beta}$ is the cross-power with the optical depth. The ensemble mean of this estimator is zero while the variance is

$$\langle \hat{\tau}_{\ell m}^{\alpha} (\hat{\tau}_{\ell m}^{\beta})^{\dagger} \rangle = \sum_{\gamma,\delta=0}^{N-1} (C^{\tau g})_{\ell}^{\alpha\gamma} ([C^{gg}]^{-1})_{\ell}^{\gamma\delta} (C^{g\tau})_{\ell}^{\delta\beta}. \qquad (41)$$

For correlated Gaussian fields Eq. (40) will yield an optical depth that has the expected statistics. However, it is likely that there are superior ways to faithfully estimate the optical depth from the galaxy density. We leave this question to future work. Note that the estimator variance reduces to $(C^{\tau\tau})_{\ell}^{\alpha\beta}$ in the limit where the actual realization of the optical depth is related to the realization of the galaxy density by $\tau_{\ell m}^{\alpha} = \sum b_{\ell}^{\alpha\gamma} g_{\ell m}^{\gamma}$; that is, when the optical depth in a redshift bin can be obtained by isotropically filtering a linear combination of the tracers in different bins.

An implementation of the MaxL estimator in this scenario involves two steps: first one estimates the optical depth field, then one uses in the MaxL estimator to obtain the reconstructed velocity. To compare the MaxL and quadratic estimators we must now include a galaxy field in our correlated Gaussian simulations. Following in the implementation in Ref. [7], we generate galaxy auto- and cross-spectra for an LSST-like galaxy survey. We refer the reader to Ref. [7] for complete details, but in brief: the galaxy spectra are generated using a halo model

incorporating a realistic halo occupation distribution (HOD), with number density of galaxies in the survey assumed to be

$$n(z) = \frac{n_g}{2z_0} \left(\frac{z}{z_0} \right)^2 \exp\left(-\frac{z}{z_0} \right), \qquad (42)$$

with $z_0 = 0.3$ and $n_g = 40$ arcmin^{-2}. The HOD and number density of objects determine the galaxy bias and galaxy shot noise in $\mathbf{C}_{\ell}^{gg}$; the halo model is used to determine a self-consistent $\mathbf{C}_{\ell}^{g\tau}$. Using these spectra as inputs, we generate properly correlated Gaussian maps of all fields at a Healpix resolution of $N_{\text{side}}^{\text{in}} = 2048$. Redshift errors associated with a photometric galaxy survey such as LSST lead to bin-bin correlations in the galaxy spectra on both large and small scales; we defer the inclusion of this effect to future work.

Using the Gaussian simulations described above, we implement the QE on the galaxy density and temperature (kSZ and primary CMB) maps using the theory spectra for $\mathbf{C}_{\ell}^{gg}$ and $\mathbf{C}_{\ell}^{g\tau}$ as inputs. We then use Eq. (40) to estimate the optical depth from the galaxy density and apply the MaxL estimator to the estimated optical depth and temperature map. The reconstruction residuals $\hat{v} - v$ for the MaxL and quadratic estimators are shown in Fig. 5. At low redshift (left panel), a high-fidelity reconstruction is possible, albeit at a lower signal-to-noise ratio than what is possible using the optical depth itself as a tracer. At high redshift all galaxy density measurements are shot noise dominated, as the galaxy density is very low. As such, it is not possible to reconstruct the high-z velocity; this is evident on the right-hand side of Fig. 5. At both low and high redshift, the MaxL estimator performs at least as well as the QE, in this regime ($N_{\text{side}}^{\text{in}} = 2048$). Given that an LSST-like experiment is shot noise dominated on small angular scales, the additional information leveraged by the MaxL estimator from the optical depth on small scales is significantly degraded. Because of this, it may be difficult to identify an
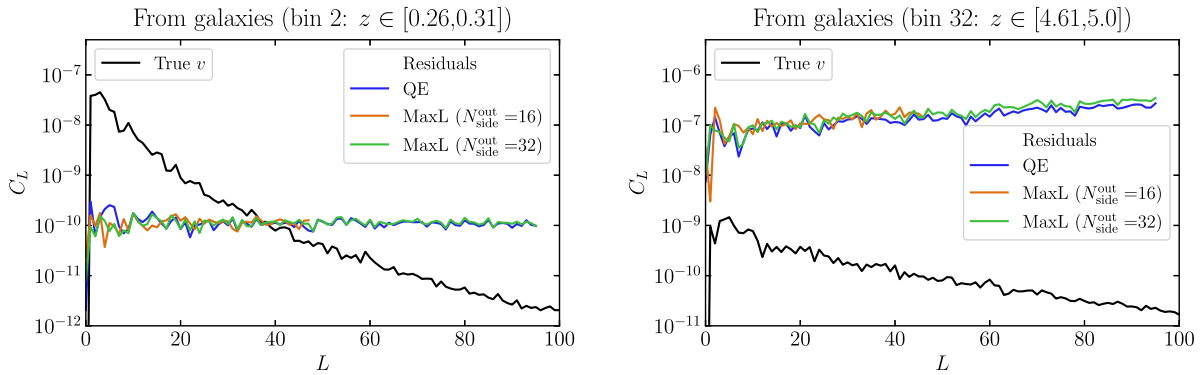

FIG. 5. The true velocity power spectrum C_L^{vv} and the power of the residuals $v - \hat{v}^{QE}$ and $v - \hat{v}^{MaxL}$, using galaxies to estimate the optical depth field. We do this for the MaxL by using the procedure defined in (40). For the QE we use the simulated galaxy maps directly and modify the theoretical weights in the estimator to ensure an unbiased output.

implementation of the MaxL estimator that would outperform the QE for near-term experiments.

When a tracer is used in the MaxL estimator, we must confront the inevitable modeling uncertainty involved in estimating the optical depth field. In the estimator Eq. (40), this modeling uncertainty arises through the theory spectra $(C^{\tau g})_\ell^{\alpha\beta}$ and $(C^{gg})_\ell^{\alpha\beta}$ that must be provided as inputs. As we demonstrate explicitly in Appendix E, the consequence is an overall multiplicative bias on the estimator mean, referred to in previous literature (e.g. [6,11–13]) as the optical depth bias. Neglecting the coarse-graining bias, for the map space estimator we have

$$\langle \hat{v}_I^\alpha \rangle = \sum_{\gamma=0}^{N-1} R_I^{\alpha\gamma} v_I^\gamma, \qquad (43)$$

where the optical depth bias matrix is

$$R_I^{\alpha\gamma} \equiv \sum_{\beta=0}^{N-1} \left[\int d\Omega\, W_I(\hat{n}) \hat{\tau}^\beta(\hat{n}) \hat{\bar{\tau}}^\alpha(\hat{n}) \right]^{-1}$$
$$\times \int d\Omega\, W_I(\hat{n}) \hat{\tau}^\beta(\hat{n}) \bar{\tau}^\gamma(\hat{n}). \qquad (44)$$

A hat on τ implies that it has been reconstructed from a tracer, while the absence of a hat indicates the actual optical depth. In complete generality, the optical depth bias matrix contains $N^2 \times N_{\text{pix}}$ free parameters. These free parameters are completely degenerate with the $N \times N_{\text{pix}}$ independent degrees of freedom in the coarse grained velocity. Fortunately, as argued in Appendix E, we expect that the optical depth bias does not depend on the pixel index I. Roughly, this follows from the assumption that the optical depth field is statistically isotropic. Furthermore, we expect that the dominant contribution to the off-diagonal $\alpha \neq \gamma$ components of the optical depth matrix are expected to arise from redshift errors in the tracer used to reconstruct the optical depth. This implies that there are only N free parameters in the optical depth bias matrix, so long as the redshift errors of the tracer can be well characterized. This is consistent with the QE, where there are N free parameters expected in the optical depth bias [7]. A complete investigation of the optical depth bias will be the focus of future work.

VII. CONCLUSIONS

In this paper, we have introduced a new class of estimators for the radial velocity field (more generally, the remote dipole field) from the CMB temperature anisotropies induced by the kinetic Sunyaev Zel'dovich effect and a tracer of the optical depth field. These MaxL estimators are based on a coarse-grained maximum likelihood condition, and we have introduced two coarse-graining schemes in harmonic and map space that can be used to solve for the maximum likelihood radial velocity

field. We derived the expected properties of the MaxL estimator, and compared the performance of the MaxL estimator to the existing quadratic estimator using the correlated Gaussian simulation framework of Ref. [7].

The MaxL estimator introduced here has several advantages over the existing QE for velocity reconstruction. As we have demonstrated, the MaxL estimator can in principle yield a higher fidelity reconstruction than the QE in the high-resolution, low-noise frontier of future CMB experiments. The MaxL estimator is not derived based on the assumption that the optical depth field is Gaussian, and it therefore does not receive contributions from the non-Gaussian bias terms present in the QE [10]. The map space version of the MaxL estimator is numerically stable and simple to implement. It is absolutely local, making it a useful and versatile tool to apply to datasets containing masks and spatially varying foregrounds. The modular nature of the MaxL estimator makes it amenable to improvements in techniques used to estimate the optical depth field or improved coarse-graining schemes. Furthermore, applying both the quadratic and MaxL estimators to future datasets will be a useful diagnostic of various biases and systematics in each analysis.

There are a variety of directions to take in future work. An important first step is a comparison of the quadratic and MaxL estimators using N-body data, where some of the advantages of the MaxL estimator described above can be confirmed in a more realistic setting. Another line of inquiry is to determine better methods for estimating the optical depth field from a galaxy redshift survey, or other tracers of large scale structure such as the cosmic infrared background [22]. For example, machine learning techniques trained on simulations could capture non-Gaussian information to better model the tracer-optical depth correlations; such a framework could incorporate multiple tracers into an overall estimate of the optical depth field over a wide range of redshifts. It will also be important to develop an analysis framework incorporating a mask and spatially varying foreground mitigation techniques as exists for the QE [7]. Finally, one can construct an analogous set of MaxL estimators based on other CMB secondaries; for example, transverse velocity reconstruction from the moving lens effect [23] or remote quadrupole reconstruction from the polarized Sunyaev Zel'dovich effect [5,24]. Improved techniques to produce unbiased high-fidelity reconstruction of cosmological fields will yield significant scientific dividends for future experiments, maximizing the ability of these datasets to constrain fundamental properties of our Universe.

APPENDIX A: HARMONIC AND MAP SPACE REPRESENTATIONS

In this paper, we work both in harmonic space and in map space (also referred to as "pixel" or "real" space). In each case, we are considering fields in a different basis: in map space, we describe fields by associating a value to each pixel on a sphere, while in harmonic space we describe fields by expanding in the basis of spherical harmonics. In this appendix, we briefly recall how to transform from one basis to another, considering both the fields themselves, and the operators on these fields.

1. Fields

A two-dimensional field $A(\hat{n})$ defined on a sphere can be represented in the map or harmonic basis:

$$a(\hat{n}) = \sum_i a_i P_i(\hat{n}) = \sum_{\ell m} a_{\ell m} Y_{\ell m}(\hat{n}), \qquad (A1)$$

where a_i is the value of the field at each pixel $P_i(\hat{n})$ and $a_{\ell m}$ are the spherical harmonic coefficients of the field decomposed onto the spherical harmonic basis $Y_{\ell m}(\hat{n})$. In the map basis, each pixel can be denoted i and the value a_i is given by

$$a_i = \int d\Omega\, a(\hat{n}) P_i(\hat{n}), \qquad (A2)$$

where P_i is an element of the discrete pixelization basis. In the harmonic basis, fields can be represented as a sum over the spherical harmonic basis $Y_{\ell m}$ with the $a_{\ell m}$'s given by

$$a_{\ell m} = \int d\Omega\, a(\hat{n}) Y^*_{\ell m}(\hat{n}). \qquad (A3)$$

In particular, we can go between map space and harmonic space according to

$$a_{\ell m} = \int d\Omega \left(\sum_i a_i P_i(\hat{n}) \right) Y^*_{\ell m}(\hat{n}) = \sum_i a_i Y^*_{\ell m}(\hat{n}_i); \quad (A4)$$

$$a_i = \int d\Omega \sum_{\ell m} a_{\ell m} Y_{\ell m}(\hat{n}) P_i(\hat{n}) = \sum_{\ell m} a_{\ell m} Y_{\ell m}(\hat{n}_i), \quad (A5)$$

where $Y_{\ell m}(\hat{n}_i) \equiv \int d\Omega P_i(\hat{n}) Y_{\ell m}(\hat{n})$, e.g. the spherical harmonic weighted over the pixel centered on $\hat{n}_i$.

2. Operators

We can decompose operators $\boldsymbol{O}$ on either basis. Thus the expression

$$\boldsymbol{O}b = c \qquad (A6)$$

can be written as

$$\sum_j \boldsymbol{O}_{ij} b_j = c_i, \qquad (A7)$$

or

$$\sum_{LM} \boldsymbol{O}_{\ell m; LM} b_{LM} = c_{\ell m}. \qquad (A8)$$

We can also consider a mixed basis

$$\sum_{LM} \boldsymbol{O}_{i; LM} b_{LM} = c_i. \qquad (A9)$$

We also have the relation

$$\boldsymbol{O}_{\ell m; LM} b_{LM} = \sum_i \boldsymbol{O}_{ij} b_j Y^*_{\ell m}(\hat{n}_i) \qquad (A10)$$

$$= \sum_i \boldsymbol{O}_{ij} \sum_{\ell' m'} b_{\ell' m'} Y_{\ell' m'}(\hat{n}_j) Y^*_{\ell m}(\hat{n}_i). \qquad (A11)$$

3. The maximum likelihood condition in map and harmonic space

The maximum likelihood condition is

$$[\boldsymbol{\tau}^\dagger (\boldsymbol{C}^{\mathrm{pCMB}})^{-1} \Theta] = [\boldsymbol{\tau}^\dagger (\boldsymbol{C}^{\mathrm{pCMB}})^{-1} \boldsymbol{\tau}] \cdot v. \qquad (A12)$$

To write this in harmonic or map space, we need to form both the operator $\boldsymbol{\tau}^\dagger (\boldsymbol{C}^{\mathrm{pCMB}})^{-1} \boldsymbol{\tau}$ and the vector $\boldsymbol{\tau}^\dagger (\boldsymbol{C}^{\mathrm{pCMB}})^{-1} \Theta$ in map and harmonic space. Let us first consider the operator $\boldsymbol{\tau}^\dagger (\boldsymbol{C}^{\mathrm{pCMB}})^{-1} \boldsymbol{\tau}$.

First, it will be useful to write the elements of $\boldsymbol{\tau}$ and $\boldsymbol{C}^{\mathrm{pCMB}}$. In map space, $\boldsymbol{\tau}$ is diagonal:

$$\boldsymbol{\tau}^\alpha_{ij} = \tau^\alpha_i \delta_{ij} \qquad (A13)$$

where τ^α_i are the components of the τ^α field in the position basis and δ_{ij} is the Kronecker delta, the identity on map space. Using the basis transformation Eq. (A11) for a diagonal operator on map space gives

$$\boldsymbol{\tau}^\alpha_{\ell m; LM} = \sum_{\ell' m'} \tau^\alpha_{\ell' m'} \int d\Omega Y_{\ell' m'}(\hat{n}) Y_{LM}(\hat{n}) Y^*_{\ell m}(\hat{n}). \qquad (A14)$$

We can also write a mixed-basis expression

$$\tau^{\alpha}_{i;LM} = \tau^{\alpha}_i Y_{LM}(\hat{n}_i). \tag{A15}$$

It will be most helpful to consider $\boldsymbol{C}^{\text{pCMB}}$ in harmonic space, where it is diagonal:

$$[(\boldsymbol{C}^{\text{pCMB}})^{-1}\boldsymbol{\tau}]^{\alpha}_{\ell''m''LM} = \sum_{\ell m}(\boldsymbol{C}^{\text{pCMB}})^{-1}_{\ell''m''\ell m}\boldsymbol{\tau}^{\alpha}_{\ell m;LM}$$

$$= \sum_{\ell m;\ell'm'}\left[\int d\Omega Y_{\ell'm'}Y_{LM}Y^*_{\ell m}\right]\frac{\tau^{\alpha}_{\ell'm'}}{C^{\text{pCMB}}_{\ell}}\delta_{\ell\ell''}\delta_{mm''}$$

$$= \sum_{\ell'm'}(-1)^{m''}\sqrt{\frac{(2\ell'+1)(2L+1)(2\ell''+1)}{4\pi}}\begin{pmatrix}L & \ell' & \ell'' \\ 0 & 0 & 0\end{pmatrix}\begin{pmatrix}L & \ell' & \ell'' \\ M & m' & -m''\end{pmatrix}\frac{\tau^{\alpha}_{\ell'm'}}{C^{\text{pCMB}}_{\ell''}}, \tag{A17}$$

where $\begin{pmatrix}L & \ell' & \ell'' \\ M & m' & -m''\end{pmatrix}$ is the Wigner $3J$ symbol. Without making further assumptions, this is as far as we can go. However, we will mainly be interested in the regime where $L \ll \ell', \ell''$. The triangle inequality restricts ℓ and ℓ' to the range $|\ell'' - \ell'| \le L \le \ell'' + \ell'$. So, as long as $C^{\text{pCMB}}_{\ell''}$ is not a strongly varying function, we can approximate:

$$C^{\text{pCMB}}_{\ell''} \simeq C^{\text{pCMB}}_{\ell''\pm L} \simeq C^{\text{pCMB}}_{\ell'}. \tag{A18}$$

Under this approximation, we have

$$\sum_{\ell m}[(\boldsymbol{C}^{\text{pCMB}})^{-1}]_{\ell''m'';\ell m}\boldsymbol{\tau}^{\alpha}_{\ell m;LM}$$

$$= \int d\Omega\left[\sum_{\ell'm'}\frac{\tau^{\alpha}_{\ell'm'}}{C^{\text{pCMB}}_{\ell'}}Y_{\ell'm'}\right]Y_{LM}Y^*_{\ell''m''}, \tag{A19}$$

which has a nice mixed pixel-space harmonic-space representation:

$$[(\boldsymbol{C}^{\text{pCMB}})^{-1}]_{ij}\boldsymbol{\tau}^{\alpha}_{j;LM} = \bar{\tau}^{\alpha}_i Y_{LM}(\hat{n}_i), \tag{A20}$$

where $\bar{\tau}^{\alpha}_i$ is the inverse-variance filtered field

$$\bar{\tau}^{\alpha}_i \equiv \sum_{\ell'm'}\frac{\tau^{\alpha}_{\ell'm'}}{C^{\text{pCMB}}_{\ell'}}Y_{\ell'm'}(\hat{n}_i). \tag{A21}$$

With these results, we can now form the operator $[\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\text{pCMB}})^{-1}\boldsymbol{\tau}]$. It is most convenient to do this using the mixed representation:

$$[\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\text{pCMB}})^{-1}\boldsymbol{\tau}]^{\alpha\beta}_{L'M';LM}$$

$$= \int d\Omega \tau^{\alpha}(\hat{n})\bar{\tau}^{\beta}(\hat{n})Y_{L'M'}(\hat{n})^*Y_{LM}(\hat{n}). \tag{A22}$$

Using an analogous set of arguments, we have

$$[\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\text{pCMB}})^{-1}\boldsymbol{\Theta}]^{\beta}_{L'M'} = \int d\Omega \tau^{\beta}(\hat{n})\bar{\Theta}(\hat{n})Y^*_{L'M'}(\hat{n}) \tag{A23}$$

where, again, the bar denotes inverse-variance filtering:

$$\bar{\Theta}_i \equiv \sum_{\ell'm'}\frac{\Theta_{\ell'm'}}{C^{\text{pCMB}}_{\ell'}}Y_{\ell'm'}(\hat{n}_i). \tag{A24}$$

The maximum likelihood condition is

$$[\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\text{pCMB}})^{-1}\boldsymbol{\Theta}]^{\alpha}_{L'M'} = \sum_{LM;\beta}[\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\text{pCMB}})^{-1}\boldsymbol{\tau}]^{\alpha\beta}_{L'M';LM}v^{\beta}_{LM}. \tag{A25}$$

We can also rotate into map space by first writing

$$[\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\text{pCMB}})^{-1}\boldsymbol{\tau}]^{\alpha\beta}_{i;LM} = \tau^{\alpha}_i\bar{\tau}^{\beta}_i Y_{LM}(\hat{n}_i) \tag{A26}$$

and

$$[\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\text{pCMB}})^{-1}\boldsymbol{\Theta}]^{\alpha}_i = \tau^{\alpha}_i\bar{\Theta}_i. \tag{A27}$$

The maximum likelihood condition is then

$$\tau^{\alpha}_i\bar{\Theta}_i = \sum_{LM;\beta}\tau^{\alpha}_i\bar{\tau}^{\beta}_i v^{\beta}_{LM}Y_{LM}(\hat{n}_i) = \sum_{\beta=0}^{N-1}\tau^{\alpha}_i\bar{\tau}^{\beta}_i v^{\beta}_i. \tag{A28}$$

APPENDIX B: COARSE-GRAINING

To determine if there is a solution to the maximum likelihood condition Eq. (A12), we must determine if the operator $\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\text{pCMB}})^{-1}\boldsymbol{\tau}$ is invertible for a particular coarse graining procedure. Note that this operator depends on the *actual realization* of the optical depth anisotropies, so there is some danger that an inverse exists for some realizations but not others. Nevertheless, we can get some idea of what

to expect by examining the statistical properties of this operator over an ensemble of realizations of the optical depth anisotropies. For example, if we focus on the harmonic basis, the mean of the operator is

$$\langle [\boldsymbol{\tau}^\dagger (\boldsymbol{C}^{\mathrm{pCMB}})^{-1} \boldsymbol{\tau}]^{\alpha\beta}_{L'M';LM} \rangle_\tau$$
$$= \sum_{\ell,\ell'} \frac{(2\ell+1)(2\ell'+1)}{4\pi} \begin{pmatrix} \ell & \ell' & L \\ 0 & 0 & 0 \end{pmatrix}^2 \frac{C^{\tau^\alpha \tau^\beta}_{\ell'}}{C^{\mathrm{pCMB}}_\ell} \delta_{LL'} \delta_{MM'}$$

$$\text{(B1)}$$

where the $\langle \ldots \rangle_\tau$ denotes an ensemble average over the optical depth anisotropies. On average, we see that the operator is diagonal in harmonic space. If the sum is dominated by large-ℓ (which we expect it to be) then it will also be nearly diagonal in bin space. This is because the bin-bin correlations in $C^{\tau^a \tau^b}_\ell$ become vanishingly small at high ℓ as the scales probed by the angular projection become much smaller than the extent of the bins. Finding the variance of the operator,

$$\langle ([\boldsymbol{\tau}^\dagger (\boldsymbol{C}^{\mathrm{pCMB}})^{-1} \boldsymbol{\tau}]^{\alpha\beta}_{L'M';LM})^2 \rangle$$
$$\simeq \langle [\boldsymbol{\tau}^\dagger (\boldsymbol{C}^{\mathrm{pCMB}})^{-1} \boldsymbol{\tau}]^{\alpha\beta}_{L'M';LM} \rangle^2$$
$$+ \left(\sum_\ell \frac{2\ell+1}{4\pi} \left[\frac{C^{\tau^\alpha \tau^\alpha}_\ell C^{\tau^\beta \tau^\beta}_\ell}{(C^{\mathrm{pCMB}}_\ell)^2} + \frac{C^{\tau^\alpha \tau^\beta}_\ell C^{\tau^\alpha \tau^\beta}_\ell}{(C^{\mathrm{pCMB}}_\ell)^2} \right] \right)$$
$$\times \sum_{\ell''} (2\ell''+1)(2L+1)(2L'+1)$$
$$\times \begin{pmatrix} L & L' & \ell'' \\ 0 & 0 & 0 \end{pmatrix}^2 \begin{pmatrix} L & L' & \ell'' \\ M & M' & M-M' \end{pmatrix}^2 \quad \text{(B2)}$$

where we have assumed that the optical depth anisotropies are Gaussian and used the limit where L, L' are small. Notably, the variance is not diagonal in the harmonic coefficients, and becomes increasingly nondiagonal at higher L and L'. It is also nondiagonal in bin space. Therefore, in a particular realization we can expect that the operator $\boldsymbol{\tau}^\dagger (\boldsymbol{C}^{\mathrm{pCMB}})^{-1} \boldsymbol{\tau}$ becomes highly nondiagonal on small angular scales. This implies that the condition number could be quite large at full resolution, and there may be no well-defined inverse. Our strategy is therefore to coarse-grain the maximum likelihood condition by truncating at some empirically determined $L_{\max}$ and N before solving for the velocity field.

APPENDIX C: CONNECTING TO THE QUADRATIC ESTIMATOR

1. Harmonic space estimator

In the low signal-to-noise regime, we can obtain the quadratic estimator from the maximum likelihood condition Eq. (21) by taking an ensemble average over the optical depth $\langle \rangle_\tau$:

$$[\boldsymbol{\tau}^\dagger (\boldsymbol{C}^{\mathrm{pCMB}})^{-1} \boldsymbol{\Theta}]^\alpha_{L'M'} = \sum_{LM;\beta} \langle [\boldsymbol{\tau}^\dagger (\boldsymbol{C}^{\mathrm{pCMB}})^{-1} \boldsymbol{\tau}]^{\alpha\beta}_{L'M';LM} \rangle_\tau v^\beta_{LM}.$$

$$\text{(C1)}$$

The ensemble average is given by Eq. (B1), which using the notation of Ref. [7] can be recognized as

$$\langle [\boldsymbol{\tau}^\dagger (\boldsymbol{C}^{\mathrm{pCMB}})^{-1} \boldsymbol{\tau}]^{\alpha\beta}_{L'M';LM} \rangle_\tau = \frac{1}{2L+1} \sum_{\ell,\ell'} G^{v^\beta \tau^\beta}_{\ell\ell'L} f^{v^\beta \tau^\alpha}_{\ell L \ell'} \delta_{LL'} \delta_{MM'}$$

$$\text{(C2)}$$

with

$$G^{v^\beta \tau^\beta}_{\ell\ell'L} = \frac{f^{v^\beta \tau^\beta}_{\ell L \ell'}}{C^{\mathrm{pCMB}}_\ell C^{\tau^\beta \tau^\beta}_{\ell'}};$$
$$f^{v^\beta \tau^\alpha}_{\ell L \ell'} = \sqrt{\frac{(2\ell+1)(2\ell'+1)(2L+1)}{4\pi}} \begin{pmatrix} \ell & \ell' & L \\ 0 & 0 & 0 \end{pmatrix} C^{\tau^\beta \tau^\alpha}_{\ell'}.$$

$$\text{(C3)}$$

The left-hand side of the maximum likelihood condition can be expanded to

$$[\boldsymbol{\tau}^\dagger (\boldsymbol{C}^{\mathrm{pCMB}})^{-1} \boldsymbol{\Theta}]^\alpha_{LM} = \sum_{\ell m;\ell'm'} (-1)^M \begin{pmatrix} \ell & \ell' & L \\ m & m' & -M \end{pmatrix}$$
$$\times G^{v^\alpha \tau^\alpha}_{\ell\ell'L} \Theta_{\ell m} \tau^\alpha_{\ell'm'}. \quad \text{(C4)}$$

For a completely diagonal $C^{\tau^\beta \tau^\alpha}_{\ell'} = C^{\tau^\alpha \tau^\alpha}_{\ell'} \delta_{\alpha\beta}$, we solve the ensemble-averaged maximum likelihood condition to arrive at the traditional form for the quadratic estimator appearing in Refs. [5,7]:

$$v^\alpha_{LM} = A^{v^\alpha}_L \sum_{\ell m;\ell'm'} (-1)^M \begin{pmatrix} \ell & \ell' & L \\ m & m' & -M \end{pmatrix} G^{v^\alpha \tau^\alpha}_{\ell\ell'L} \Theta_{\ell m} \tau^\alpha_{\ell'm'}$$

$$\text{(C5)}$$

where

$$A^{v^\alpha}_L \equiv \left[\frac{1}{2L+1} \sum_{\ell,\ell'} G^{v^\alpha \tau^\alpha}_{\ell\ell'L} f^{v^\alpha \tau^\alpha}_{\ell L \ell'} \right]^{-1}. \quad \text{(C6)}$$

There are, however, a few crucial differences with the quadratic estimator presented in Refs. [5,7]. In the traditional quadratic estimator, the weights needed for an unbiased minimum variance reconstruction require that everywhere above we replace $C^{\mathrm{pCMB}}_\ell \to C^{\mathrm{pCMB}}_\ell + C^{\mathrm{kSZ}}_\ell$. In the limit where $C^{\mathrm{pCMB}}_\ell \ll C^{\mathrm{kSZ}}_\ell$, as occurs for a sufficiently high-resolution low-noise CMB experiment, this is a non-negligible difference—as expected, since in this case the temperature anisotropies on small scales are very non-Gaussian, and far from the limit of Gaussian fields used to

derive the quadratic estimator. Therefore, we only recover the traditional quadratic estimator in the limit where $C_\ell^{\text{pCMB}} \gg C_\ell^{\text{kSZ}}$. This observation supports the notion that the maximum likelihood estimator will be superior in the high signal-to-noise regime. Another difference arises in the case where $C_\ell^{\tau^\beta \tau^\alpha}$ is not diagonal. In this case, to find a solution to the ensemble averaged maximum likelihood condition it is necessary to find the inverse of $C_\ell^{\tau^\beta \tau^\alpha}$ (if it exists). The resulting solution is different than the approach presented in Ref. [7], which first constructs a set of biased estimators which are then de-biased by forming the appropriate linear combinations. The relevant differences would be interesting to explore further in future work.

2. Map space estimator

Just as for the harmonic space estimator, we can derive an ensemble average map space estimator. Starting from Eq. (32):

$$\int d\Omega \, W_I(\hat{n}) \tau^\alpha(\hat{n}) \bar{\Theta}(\hat{n}) = \sum_{\beta=0}^{N-1} v_I^\beta \int d\Omega \, W_I(\hat{n}) \langle \tau^\alpha(\hat{n}) \bar{\tau}^\beta(\hat{n}) \rangle_\tau \tag{C7}$$

$$+ \sum_{\beta=0}^{N-1} \int d\Omega \, W_I(\hat{n}) \langle \tau^\alpha(\hat{n}) \bar{\tau}^\beta(\hat{n}) \rangle_\tau \Delta v_I^\beta(\hat{n}) \tag{C8}$$

$$= \sum_{\beta=0}^{N-1} v_I^\beta \langle \tau^\alpha(0) \bar{\tau}^\beta(0) \rangle_\tau \int d\Omega \left[W_I(\hat{n}) + \int d\Omega W_I(\hat{n}) \Delta v_I^\beta(\hat{n}) \right] \tag{C9}$$

$$= \sum_{\beta=0}^{N-1} v_I^\beta \langle \tau^\alpha(0) \bar{\tau}^\beta(0) \rangle_\tau \tag{C10}$$

where we have assumed that τ^α is statistically isotropic, $\int d\Omega \, W_I(\hat{n}) = 1$, and $\int d\Omega \, W_I(\hat{n}) \Delta v_I^\beta(\hat{n}) = 0$ by definition. This suggests the following map space quadratic estimator:

$$\hat{v}_I^\alpha = \sum_{\beta=0}^{N-1} [\langle \tau^\beta(0) \bar{\tau}^\alpha(0) \rangle_\tau]^{-1} \int d\Omega \, W_I(\hat{n}) \tau^\alpha(\hat{n}) \bar{\Theta}(\hat{n}). \tag{C11}$$

For a diagonal variance $\langle \tau^\beta(0) \bar{\tau}^\alpha(0) \rangle_\tau = \langle \tau^\alpha(0) \bar{\tau}^\alpha(0) \rangle_\tau \delta_{\alpha\beta}$, we can identify

$$\langle \tau^\alpha(0) \bar{\tau}^\alpha(0) \rangle_\tau = \frac{1}{A_{L=0}^{v^\alpha}} \tag{C12}$$

as defined above in Eq. (C6). Since Eq. (C6) is nearly independent of L [5,7], one can simply show that Eq. (C11) is equivalent to Eq. (C5) at low L, and also equivalent to the map based quadratic estimator presented in Refs. [5,7] (with the same caveats as for the harmonic space estimator). In the low signal-to-noise regime, Eq. (C11) may be a simpler implementation of the quadratic estimator used in previous work. We leave further investigation of this point to future work.

APPENDIX D: MEAN AND VARIANCE OF THE MAXL ESTIMATOR

1. Harmonic space estimator

The mean of the harmonic space estimator is

$$\langle \hat{v}_{LM}^\alpha \rangle_\Theta = \sum_{L'M'}^{L'_{\text{max}}} \sum_{\beta=0}^{N-1} ([\boldsymbol{\tau}^\dagger (C^{\text{pCMB}})^{-1} \boldsymbol{\tau}]^{-1})_{LM;L'M'}^{\alpha\beta}$$
$$\times \langle [\boldsymbol{\tau}^\dagger (C^{\text{pCMB}})^{-1} \Theta]_{L'M'}^\beta \rangle_\Theta. \tag{D1}$$

To calculate this, we must evaluate

$$\langle [\boldsymbol{\tau}^\dagger (C^{\text{pCMB}})^{-1} \Theta]_{L'M'}^\alpha \rangle_\Theta = \int d\Omega \, \tau^\alpha(\hat{n}) \langle \bar{\Theta}(\hat{n}) \rangle_\Theta Y_{L'M'}^*(\hat{n}). \tag{D2}$$

The observed temperature anisotropies consist of the zero-mean primary CMB and the kSZ:

$$\langle \bar{\Theta}(\hat{n}) \rangle_\Theta = \langle \bar{\Theta}^{\text{pCMB}}(\hat{n}) \rangle_\Theta + \bar{\Theta}^{\text{kSZ}}(\hat{n}) = \bar{\Theta}^{\text{kSZ}}(\hat{n}). \tag{D3}$$

Expanding the kSZ anisotropies into coarse- and fine-grained components, we have

$$\langle[\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\mathrm{pCMB}})^{-1}\Theta]^{\alpha}_{L'M'}\rangle_{\Theta} = \int d\Omega\tau^{\alpha}(\hat{n})\Bigg[\sum_{\beta=0}^{N-1}\sum_{LM}^{L_{\max}}\bar{\tau}^{\beta}(\hat{n})v^{\beta}_{LM}Y_{LM}(\hat{n})$$

$$+ \sum_{b=N}^{\infty}\sum_{LM}^{L_{\max}}\bar{\tau}^{b}(\hat{n})v^{b}_{LM}Y_{LM}(\hat{n}) + \sum_{b=0}^{\infty}\sum_{LM=L_{\max}}^{\infty}\bar{\tau}^{b}(\hat{n})v^{b}_{LM}Y_{LM}(\hat{n})\Bigg]Y^{*}_{L'M'}(\hat{n})$$

$$\simeq \sum_{\beta=0}^{N-1}\sum_{LM}^{L_{\max}}[\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\mathrm{pCMB}})^{-1}\boldsymbol{\tau}]^{\alpha\beta}_{L'M';LM}v^{\beta}_{LM} + \sum_{b=N}^{\infty}\sum_{LM}^{L_{\max}}[\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\mathrm{pCMB}})^{-1}\boldsymbol{\tau}]^{\alpha b}_{L'M';LM}v^{b}_{LM}$$

where we assumed that the dominant contributions from the velocity field are at low L. We therefore have

$$\langle\hat{v}^{\alpha}_{LM}\rangle_{\Theta} = v^{\alpha}_{LM} + \beta^{\alpha}_{LM} \tag{D4}$$

where

$$\beta^{\alpha}_{LM} \equiv \sum_{L'M'}^{L_{\max}}\sum_{\gamma=0}^{N-1}\sum_{c=N}^{\infty}\sum_{L''M''}^{L_{\max}}([\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\mathrm{pCMB}})^{-1}\boldsymbol{\tau}]^{-1})^{\alpha\gamma}_{LM;L'M'}[\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\mathrm{pCMB}})^{-1}\boldsymbol{\tau}]^{\gamma c}_{L'M';L''M''}v^{c}_{L''M''}. \tag{D5}$$

So, the estimator is unbiased up to an additive bias β^{α}_{LM}, which depends on the fine-grained radial velocity modes.

Moving to the variance,

$$\langle\hat{v}^{\alpha}_{LM}\hat{v}^{\beta}_{L'M'}\rangle_{\Theta} = \sum_{L_aM_a;L_bM_b}^{L_{\max}}\sum_{\gamma;\delta=0}^{N-1}([\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\mathrm{pCMB}})^{-1}\boldsymbol{\tau}]^{-1})^{\alpha\gamma}_{LM;L_aM_a}([\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\mathrm{pCMB}})^{-1}\boldsymbol{\tau}]^{-1})^{\beta\delta}_{L'M';L_bM_b}$$

$$\times \iint d\Omega d\Omega'\tau^{\gamma}(\hat{n})\tau^{\delta}(\hat{n}')\langle\bar{\Theta}(\hat{n})\bar{\Theta}(\hat{n}')\rangle_{\Theta}Y_{L_aM_a}(\hat{n})Y_{L_bM_b}(\hat{n}'). \tag{D6}$$

There is a contribution to the temperature correlation function from the primary CMB and from the kSZ effect:

$$\langle\bar{\Theta}(\hat{n})\bar{\Theta}(\hat{n}')\rangle_{\Theta} = \sum_{\ell m}\frac{1}{C^{\mathrm{pCMB}}_{\ell}}Y_{\ell m}(\hat{n})Y_{\ell m}(\hat{n}') + \bar{\Theta}^{\mathrm{kSZ}}(\hat{n})\bar{\Theta}^{\mathrm{kSZ}}(\hat{n}'). \tag{D7}$$

The first term will give rise to the Gaussian estimator noise and the second will describe our signal and biases arising due to our coarse-graining of the maximum likelihood condition:

$$\langle\hat{v}^{\alpha}_{LM}\hat{v}^{\beta}_{L'M'}\rangle_{\Theta} = (C^{vv})^{\alpha\beta}_{LM;L'M'} + N^{\alpha\beta}_{LM;L'M'}. \tag{D8}$$

The reconstruction noise $N^{\alpha\beta}_{LM;L'M'}$ arises due to the contribution from the primary CMB:

$$N^{\alpha\beta}_{LM;L'M'} = \sum_{L_aM_a;L_bM_b;\ell m}\sum_{\gamma;\delta}([\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\mathrm{pCMB}})^{-1}\boldsymbol{\tau}]^{-1})^{\alpha\gamma}_{LM;L_aM_a}([\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\mathrm{pCMB}})^{-1}\boldsymbol{\tau}]^{-1})^{\beta\delta}_{L'M';L_bM_b}$$

$$\times \iint d\Omega d\Omega'\tau^{\gamma}(\hat{n})\tau^{\delta}(\hat{n}')\frac{1}{C^{\mathrm{pCMB}}_{\ell}}Y_{\ell m}(\hat{n})Y_{\ell m}(\hat{n}')Y_{L_aM_a}(\hat{n})Y_{L_bM_b}(\hat{n}') \tag{D9}$$

$$\simeq \sum_{L_aM_a;L_bM_b}\sum_{\gamma;\delta}([\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\mathrm{pCMB}})^{-1}\boldsymbol{\tau}]^{-1})^{\alpha\gamma}_{LM;L_aM_a}([\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\mathrm{pCMB}})^{-1}\boldsymbol{\tau}]^{-1})^{\beta\delta}_{L'M';L_bM_b}\int d\Omega\tau^{\gamma}(\hat{n})\bar{\tau}^{\delta}(\hat{n})Y_{L_aM_a}(\hat{n})Y_{L_bM_b}(\hat{n})$$

$$\tag{D10}$$

$$= \sum_{L_aM_a;L_bM_b}\sum_{\gamma;\delta}([\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\mathrm{pCMB}})^{-1}\boldsymbol{\tau}]^{-1})^{\alpha\gamma}_{LM;L_aM_a}([\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\mathrm{pCMB}})^{-1}\boldsymbol{\tau}]^{-1})^{\beta\delta}_{L'M';L_bM_b} \tag{D11}$$

$$\times(\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\mathrm{pCMB}})^{-1}\boldsymbol{\tau})^{\gamma\delta}_{L_bM_b;L_aM_a} \tag{D12}$$

$$= ([\boldsymbol{\tau}^{\dagger}(\boldsymbol{C}^{\mathrm{pCMB}})^{-1}\boldsymbol{\tau}]^{-1})^{\alpha\beta}_{LM;L'M'} \tag{D13}$$

where in going from the first to second lines we have assumed that C_ℓ^{pCMB} does not vary too quickly on scales of order $\Delta\ell = L$ and used the identity

$$\int d^2 n \, Y_{\ell_a m_a} Y_{\ell_b m_b} Y_{\ell_c m_c} Y_{\ell_d m_d} = \sum_{\ell m} \left(\int d^2 n Y_{\ell_a m_a} Y_{\ell_b m_b} Y_{\ell m} \right) \left(\int d^2 n' Y_{\ell_c m_c} Y_{\ell_d m_d} Y_{\ell m} \right). \tag{D14}$$

For the signal and bias due to "fine modes" we have

$$(C^{vv})_{LM;L'M'}^{\alpha\beta} = \sum_{L_a M_a; L_b M_b} \sum_{\gamma;\delta} ([\boldsymbol{\tau}^\dagger (\boldsymbol{C}^{\mathrm{pCMB}})^{-1} \boldsymbol{\tau}]^{-1})_{LM;L_a M_a}^{\alpha\gamma} ([\boldsymbol{\tau}^\dagger (\boldsymbol{C}^{\mathrm{pCMB}})^{-1} \boldsymbol{\tau}]^{-1})_{L'M';L_b M_b}^{\beta\delta}$$
$$\times \iint d\Omega d\Omega' \tau^\gamma(\hat{n}) \tau^\delta(\hat{n}') \bar{\Theta}^{\mathrm{kSZ}}(\hat{n}) \bar{\Theta}^{\mathrm{kSZ}}(\hat{n}') Y_{L_a M_a}(\hat{n}) Y_{L_b M_b}(\hat{n}') \tag{D15}$$

$$\simeq \sum_{L_a M_a; L_b M_b} \sum_{\gamma;\delta;e;f} ([\boldsymbol{\tau}^\dagger (\boldsymbol{C}^{\mathrm{pCMB}})^{-1} \boldsymbol{\tau}]^{-1})_{LM;L_a M_a}^{\alpha\gamma} ([\boldsymbol{\tau}^\dagger (\boldsymbol{C}^{\mathrm{pCMB}})^{-1} \boldsymbol{\tau}]^{-1})_{L'M';L_b M_b}^{\beta\delta}$$
$$\times \iint d\Omega d\Omega' \tau^\gamma(\hat{n}) \tau^\delta(\hat{n}') \bar{\tau}^e(\hat{n}) v^e(\hat{n}) \bar{\tau}^f(\hat{n}') v^f(\hat{n}') Y_{L_a M_a}(\hat{n}) Y_{L_b M_b}(\hat{n}')$$
$$= \sum_{L_a M_a; L_b M_b} \sum_{\gamma;\delta;e;f} ([\boldsymbol{\tau}^\dagger (\boldsymbol{C}^{\mathrm{pCMB}})^{-1} \boldsymbol{\tau}]^{-1})_{LM;L_a M_a}^{\alpha\gamma} ([\boldsymbol{\tau}^\dagger (\boldsymbol{C}^{\mathrm{pCMB}})^{-1} \boldsymbol{\tau}]^{-1})_{L'M';L_b M_b}^{\beta\delta}$$
$$\times \sum_{L_c M_c; L_d M_d} v_{L_c M_c}^e v_{L_d M_d}^f \int d\Omega \tau^\gamma(\hat{n}) \bar{\tau}^e(\hat{n}) Y_{L_a M_a}(\hat{n}) Y_{L_c M_c}(\hat{n}) \int d\Omega' \tau^\delta(\hat{n}') \bar{\tau}^f(\hat{n}') Y_{L_b M_b}(\hat{n}') Y_{L_d M_d}(\hat{n}')$$
$$= \sum_{L_a M_a; L_b M_b} \sum_{L_c M_c; L_d M_d} \sum_{\gamma;\delta;e;f} v_{L_c M_c}^e v_{L_d M_d}^f ([\boldsymbol{\tau}^\dagger (\boldsymbol{C}^{\mathrm{pCMB}})^{-1} \boldsymbol{\tau}]^{-1})_{LM;L_a M_a}^{\alpha\gamma} [\boldsymbol{\tau}^\dagger (\boldsymbol{C}^{\mathrm{pCMB}})^{-1} \boldsymbol{\tau}]_{L_a M_a; L_c M_c}^{\gamma e}$$
$$\times ([\boldsymbol{\tau}^\dagger (\boldsymbol{C}^{\mathrm{pCMB}})^{-1} \boldsymbol{\tau}]^{-1})_{L'M';L_b M_b}^{\beta\delta} [\boldsymbol{\tau}^\dagger (\boldsymbol{C}^{\mathrm{pCMB}})^{-1} \boldsymbol{\tau}]_{L_b M_b; L_d M_d}^{\delta f}$$
$$= (v_{LM}^\alpha + \beta_{LM}^\alpha)(v_{L'M'}^\beta + \beta_{L'M'}^\beta). \tag{D16}$$

The main assumption made in the derivation of the estimator mean and variance was that C_ℓ^{pCMB} does not vary too quickly on scales of order $\Delta\ell = L$, which can always be satisfied for sufficiently small L. For modest values of $L \sim 100$, we expect this to be a reasonable assumption at large ℓ where the primary CMB is smooth and dominated by lensing, instrumental noise, or reionization kSZ (depending on the instrument).

2. Map space estimator

We now compute the mean and variance of the map space estimator. Starting with the mean:

$$\langle \hat{v}_I^\alpha \rangle_\Theta = \sum_{\beta=0}^{N-1} \boldsymbol{\mathcal{W}}_{I;\alpha\beta}^{-1} \int d\Omega \, W_I(\hat{n}) \tau^\beta(\hat{n}) \bar{\Theta}^{\mathrm{kSZ}}(\hat{n}), \tag{D17}$$

$$= \sum_{\beta=0}^{N-1} \boldsymbol{\mathcal{W}}_{I;\alpha\beta}^{-1} \int d\Omega \, W_I(\hat{n}) \tau^\beta(\hat{n}) \sum_{c=0}^{\infty} \bar{\tau}^c(\hat{n})(v_I^c + \Delta v_I^c(\hat{n})) \tag{D18}$$

$$= \sum_{\beta;\gamma=0}^{N-1} v_I^\gamma \boldsymbol{\mathcal{W}}_{I;\alpha\beta}^{-1} \boldsymbol{\mathcal{W}}_{I;\gamma\beta} + \sum_{\beta;\gamma=0}^{N-1} \boldsymbol{\mathcal{W}}_{I;\alpha\beta}^{-1} \int d\Omega \, W_I(\hat{n}) \tau^\beta(\hat{n}) \bar{\tau}^\gamma(\hat{n}) \Delta v_I^\gamma(\hat{n}) \tag{D19}$$

$$+ \sum_{\beta=0}^{N-1} \sum_{c=N}^{\infty} \boldsymbol{\mathcal{W}}_{I;\alpha\beta}^{-1} \int d\Omega \, W_I(\hat{n}) \tau^\beta(\hat{n}) \bar{\tau}^c(\hat{n}) v^c(\hat{n}) \tag{D20}$$

$$= v_I^\alpha + \beta_I^\alpha \tag{D21}$$

where we have defined

$$\beta_I^\alpha = \sum_{\beta;\gamma=0}^{N-1} \mathcal{W}_{I;\alpha\beta}^{-1} \int d\Omega\, W_I(\hat{n})\tau^\beta(\hat{n})\bar{\tau}^\gamma(\hat{n})\Delta v_I^\gamma(\hat{n}) + \sum_{\beta=0}^{N-1}\sum_{c=N}^\infty \mathcal{W}_{I;\alpha\beta}^{-1} \int d\Omega\, W_I(\hat{n})\tau^\beta(\hat{n})\bar{\tau}^c(\hat{n})v^c(\hat{n}). \qquad \text{(D22)}$$

Moving to the estimator variance, we have

$$\langle \hat{v}_I^\alpha \hat{v}_J^\beta \rangle_\Theta = \sum_{\gamma;\delta=0}^{N-1} \mathcal{W}_{I;\alpha\gamma}^{-1}\mathcal{W}_{J;\beta\delta}^{-1} \int d\Omega d\Omega' W_I(\hat{n})W_J(\hat{n}')\tau^\gamma(\hat{n})\tau^\delta(\hat{n}')\langle \bar{\Theta}(\hat{n})\bar{\Theta}(\hat{n}')\rangle_\Theta$$

$$= \sum_{\gamma;\delta=0}^{N-1} \mathcal{W}_{I;\alpha\gamma}^{-1}\mathcal{W}_{J;\beta\delta}^{-1}\left[\int d\Omega\, W_I(\hat{n})\tau^\gamma(\hat{n})\bar{\Theta}^{\text{kSZ}}(\hat{n}) \int d\Omega' W_J(\hat{n}')\tau^\delta(\hat{n}')\bar{\Theta}^{\text{kSZ}}(\hat{n}')\right.$$

$$\left. + \sum_{\ell m} \int d\Omega\, W_I(\hat{n})\tau^\gamma(\hat{n})Y_{\ell m}(\hat{n}) \int d\Omega'\, W_J(\hat{n}')\frac{\tau^\delta(\hat{n}')}{C_\ell^{\text{pCMB}}}Y_{\ell m}(\hat{n}')\right]$$

$$\simeq (v_I^\alpha + \beta_I^\alpha)(v_J^\beta + \beta_J^\beta) + \sum_{\gamma;\delta=0}^{N-1} \mathcal{W}_{I;\alpha\gamma}^{-1}\mathcal{W}_{J;\beta\delta}^{-1} \int d\Omega W_I(\hat{n})W_J(\hat{n})\tau^\gamma(\hat{n})\bar{\tau}^\delta(\hat{n})$$

$$= (v_I^\alpha + \beta_I^\alpha)(v_J^\beta + \beta_J^\beta) + \frac{1}{\Delta\theta^2}\mathcal{W}_{I;\alpha\beta}^{-1}\delta_{IJ}, \qquad \text{(D23)}$$

where in going from the second to the third line, we have made the approximation that C_ℓ^{pCMB} is not too rapidly varying with scale, as we did in the derivation of the harmonic space estimator variance.

APPENDIX E: OPTICAL DEPTH BIAS

When we only have access to an estimate $\hat{\tau}$ of the differential optical depth, the harmonic space estimator is given by

$$\hat{v}_{LM}^\alpha = \sum_{L'M'}^{L'_{\max}}\sum_{\beta=0}^{N-1} ([\hat{\boldsymbol{\tau}}^\dagger(\boldsymbol{C}^{\text{pCMB}})^{-1}\hat{\boldsymbol{\tau}}]^{-1})_{LM;L'M'}^{\alpha\beta}$$

$$\times (\hat{\boldsymbol{\tau}}^\dagger(\boldsymbol{C}^{\text{pCMB}})^{-1}\Theta)_{L'M'}^\beta. \qquad \text{(E1)}$$

Computing the estimator mean, we have

$$\langle \hat{v}_{LM}^\alpha \rangle_\Theta = \sum_{L'M'}^{L'_{\max}}\sum_{\beta=0}^{N-1} ([\hat{\boldsymbol{\tau}}^\dagger(\boldsymbol{C}^{\text{pCMB}})^{-1}\hat{\boldsymbol{\tau}}]^{-1})_{LM;L'M'}^{\alpha\beta}$$

$$\times \langle(\hat{\boldsymbol{\tau}}^\dagger(\boldsymbol{C}^{\text{pCMB}})^{-1}\Theta)_{L'M'}^\beta\rangle_\Theta \qquad \text{(E2)}$$

$$= \left[\sum_{L''M''}^{L_{\max}}\sum_{\gamma=0}^{N-1} R_{LM;L''M''}^{\alpha\gamma} v_{L''M''}^\gamma\right] + \hat{\beta}_{LM}^\alpha, \qquad \text{(E3)}$$

where

$$R_{LM;L''M''}^{\alpha\gamma} = \sum_{L'M'}^{L'_{\max}}\sum_{\beta=0}^{N-1} ([\hat{\boldsymbol{\tau}}^\dagger(\boldsymbol{C}^{\text{pCMB}})^{-1}\hat{\boldsymbol{\tau}}]^{-1})_{LM;L'M'}^{\alpha\beta}$$

$$\times [\hat{\boldsymbol{\tau}}^\dagger(\boldsymbol{C}^{\text{pCMB}})^{-1}\boldsymbol{\tau}]_{L'M';L''M''}^{\beta\gamma} \qquad \text{(E4)}$$

and

$$\hat{\beta}_{LM}^\alpha \equiv \sum_{L'M'}^{L_{\max}}\sum_{\gamma=0}^{N-1}\sum_{c=N}^\infty\sum_{L''M''}^{L_{\max}} ([\hat{\boldsymbol{\tau}}^\dagger(\boldsymbol{C}^{\text{pCMB}})^{-1}\hat{\boldsymbol{\tau}}]^{-1})_{LM;L'M'}^{\alpha\gamma}$$

$$\times [\hat{\boldsymbol{\tau}}^\dagger(\boldsymbol{C}^{\text{pCMB}})^{-1}\boldsymbol{\tau}]_{L'M';L''M''}^{\gamma c} v_{L''M''}^c. \qquad \text{(E5)}$$

When $\hat{\tau} = \tau$, this is simply the identity matrix, and we recover an unbiased estimator mean. In reality, an imperfect estimate of the differential optical depth will result in the mixing of information and therefore a bias. It is important to characterize the magnitude and properties of this bias to obtain useful cosmological information from the reconstruction.

To explore the optical depth bias a bit more, suppose we have a pretty good reconstruction of the differential optical depth, such that

$$\hat{\tau}^\alpha(\hat{n}) = \tau^\alpha(\hat{n}) + n_\tau^\alpha(\hat{n}); \qquad n_\tau^\alpha(\hat{n}) \ll \tau^\alpha(\hat{n}). \qquad \text{(E6)}$$

We can approximate

$$[\hat{\boldsymbol{\tau}}^\dagger(\boldsymbol{C}^{\text{pCMB}})^{-1}\hat{\boldsymbol{\tau}}]^{-1}$$

$$\simeq [1 - [\boldsymbol{\tau}^\dagger(\boldsymbol{C}^{\text{pCMB}})^{-1}\boldsymbol{\tau}]^{-1}\cdot[\boldsymbol{n}_\tau^\dagger(\boldsymbol{C}^{\text{pCMB}})^{-1}\boldsymbol{n}_\tau]]$$

$$\cdot [\boldsymbol{\tau}^\dagger(\boldsymbol{C}^{\text{pCMB}})^{-1}\boldsymbol{\tau}]^{-1} \qquad \text{(E7)}$$

and therefore

$$
R^{\alpha\gamma}_{LM;L''M''} \simeq \delta_{\alpha\gamma}\delta_{LM;L''M''}
$$
$$
- \sum_{L'M'}^{L'_{max}}\sum_{\beta=0}^{N-1}([\boldsymbol{\tau}^\dagger(\boldsymbol{C}^{pCMB})^{-1}\boldsymbol{\tau}]^{-1})^{\alpha\beta}_{LM;L'M'}
$$
$$
\times [\boldsymbol{n_\tau}^\dagger(\boldsymbol{C}^{pCMB})^{-1}\boldsymbol{n_\tau}]^{\beta\gamma}_{L'M';L''M''}. \quad \text{(E8)}
$$

Taking an ensemble average over τ and n_τ, and using Eq. (B1), we have

$$
\langle R^{\alpha\gamma}_{LM;L''M''}\rangle \simeq \delta_{\alpha\gamma}\delta_{LM;L''M''}
$$
$$
- \delta_{LM;L''M''}\sum_{\beta}^{\beta_{max}}\langle\tau^\alpha(0)\bar{\tau}^\beta(0)\rangle^{-1}\langle n_\tau^\beta(0)\bar{n}_\tau^\gamma(0)\rangle
$$
$$
\text{(E9)}
$$

where $\langle\tau^\alpha(0)\bar{\tau}^\beta(0)\rangle^{-1}$ is the bin-space inverse of Eq. (B1) and $\langle n_\tau^\beta(0)\bar{n}_\tau^\gamma(0)\rangle$ is

$$
\langle n_\tau^\beta(0)\bar{n}_\tau^\gamma(0)\rangle = \sum_\ell \frac{2\ell+1}{4\pi}\frac{C_\ell^{n_\tau^\beta n_\tau^\gamma}}{C_\ell^{pCMB}}. \quad \text{(E10)}
$$

In particular, note that on average in this limit the optical depth bias is independent of scale $L''M''$, and small when the differential optical depth is reconstructed well.

For the map space estimator, we have

$$
\hat{v}_I^\alpha = \sum_{\beta=0}^{N-1}\left[\int d\Omega\, W_I(\hat{n})\hat{\tau}^\beta(\hat{n})\hat{\bar{\tau}}^\alpha(\hat{n})\right]^{-1}
$$
$$
\times \int d\Omega\, W_I(\hat{n})\hat{\tau}^\beta(\hat{n})\bar{\Theta}(\hat{n}). \quad \text{(E11)}
$$

Computing the estimator mean,

$$
\langle\hat{v}_I^\alpha\rangle = \sum_{\gamma=0}^{N-1}R_I^{\alpha\gamma}v_I^\gamma + \hat{\beta}_I^\alpha \quad \text{(E12)}
$$

where

$$
R_I^{\alpha\gamma} \equiv \sum_{\beta=0}^{N-1}\left[\int d\Omega\, W_I(\hat{n})\hat{\tau}^\beta(\hat{n})\hat{\bar{\tau}}^\alpha(\hat{n})\right]^{-1}
$$
$$
\times \int d\Omega\, W_I(\hat{n})\hat{\tau}^\beta(\hat{n})\bar{\tau}^\gamma(\hat{n}) \quad \text{(E13)}
$$

and

$$
\hat{\beta}_I^\alpha = \sum_{\beta;\gamma=0}^{N-1}\left[\int d\Omega\, W_I(\hat{n})\hat{\tau}^\beta(\hat{n})\hat{\bar{\tau}}^\alpha(\hat{n})\right]^{-1}
$$
$$
\times \int d\Omega\, W_I(\hat{n})\hat{\tau}^\beta(\hat{n})\bar{\tau}^\gamma(\hat{n})\Delta v_I^\gamma(\hat{n})
$$
$$
+ \sum_{\beta=0}^{N-1}\sum_{c=N}^{\infty}\left[\int d\Omega\, W_I(\hat{n})\hat{\tau}^\beta(\hat{n})\hat{\bar{\tau}}^\alpha(\hat{n})\right]^{-1}
$$
$$
\times \int d\Omega\, W_I(\hat{n})\tau^\beta(\hat{n})\bar{\tau}^c(\hat{n})v^c(\hat{n}). \quad \text{(E14)}
$$

Assuming a well-modeled optical depth, as in Eq. (E6), a similar derivation as for the harmonic estimator yields

$$
\langle R_I^{\alpha\gamma}\rangle \simeq \delta_{\alpha\gamma} - \sum_{\beta=0}^{N-1}\langle\tau^\alpha(0)\bar{\tau}^\beta(0)\rangle^{-1}\langle n_\tau^\beta(0)\bar{n}_\tau^\gamma(0)\rangle. \quad \text{(E15)}
$$

So, the optical depth bias in this limit is independent of the pixel index I.

---

[1] J. Aguirre *et al.* (Simons Observatory Collaboration), J. Cosmol. Astropart. Phys. 02 (**2019**) 056.

[2] K. N. Abazajian *et al.* (CMB-S4 Collaboration), arXiv: 1610.02743.

[3] N. Sehgal *et al.*, Bull. Am. Astron. Soc. **51**, 6 (2019), https://ui.adsabs.harvard.edu/abs/2019BAAS...51g...6S/abstract.

[4] R. A. Sunyaev and I. B. Zeldovich, Mon. Not. R. Astron. Soc. **190**, 413 (1980).

[5] A.-S. Deutsch, E. Dimastrogiovanni, M. C. Johnson, M. Münchmeyer, and A. Terrana, Phys. Rev. D **98**, 123501 (2018).

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

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

[8] A. Terrana, M.-J. Harris, and M. C. Johnson, J. Cosmol. Astropart. Phys. 02 (2017) 040.

[9] J. I. Cayuso, M. C. Johnson, and J. B. Mertens, Phys. Rev. D **98**, 063502 (2018).

[10] U. Giri and K. M. Smith, J. Cosmol. Astropart. Phys. 09 (2022) 028.

[11] M. S. Madhavacheril, N. Battaglia, K. M. Smith, and J. L. Sievers, Phys. Rev. D **100**, 103532 (2019).

[12] N. Battaglia, J. Cosmol. Astropart. Phys. 08 (2016) 058.

[13] S. C. Hotinli, K. M. Smith, M. S. Madhavacheril, and M. Kamionkowski, Phys. Rev. D **104**, 083529 (2021).

[14] L. Thiele, F. Villaescusa-Navarro, D. N. Spergel, D. Nelson, and A. Pillepich, Astrophys. J. **902**, 129 (2020).

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

[16] S. C. Hotinli and M. C. Johnson, Phys. Rev. D **105**, 063522 (2022).

[17] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann, Astrophys. J. **622**, 759 (2005).

[18] Y. Akrami *et al.* (Planck Collaboration), Astron. Astrophys. **644**, A100 (2020).

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

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

[21] N. Battaglia, J. Cosmol. Astropart. Phys. 08 (2016) 058.

[22] F. McCarthy and M. C. Johnson, Phys. Rev. D **102**, 043520 (2020).

[23] S. C. Hotinli, J. Meyers, N. Dalal, A. H. Jaffe, M. C. Johnson, J. B. Mertens, M. Münchmeyer, K. M. Smith, and A. van Engelen, Phys. Rev. Lett. **123**, 061301 (2019).

[24] A.-S. Deutsch, M. C. Johnson, M. Münchmeyer, and A. Terrana, J. Cosmol. Astropart. Phys. 04 (2018) 034.