

Field-level multiprobe analysis of the CMB, integrated Sachs-Wolfe effect, and the galaxy density maps

Alan Junzhe Zhou¹ and Scott Dodelson¹

*Department of Physics, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, USA;
McWilliams Center for Cosmology, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, USA
and NSF AI Planning Institute, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, USA*

 (Received 3 April 2023; accepted 12 September 2023; published 3 October 2023)

Extracting information from cosmic surveys is often done in a two-step process, construction of maps and then summary statistics such as two-point functions. We use simulations to demonstrate the advantages of a general Bayesian framework that consistently combines different cosmological experiments on the field level, and reconstructs both the maps and cosmological parameters. We apply our method to jointly reconstruct the primordial CMB, the integrated Sachs-Wolfe effect, and six tomographic galaxy density maps on the full sky on large scales along with several cosmological parameters. While the traditional maximum a posterior estimator has both two-point level and field-level bias, the new approach yields unbiased cosmological constraints and improves the signal-to-noise ratio of the maps.

DOI: [10.1103/PhysRevD.108.083506](https://doi.org/10.1103/PhysRevD.108.083506)

I. INTRODUCTION

The large-scale structure (LSS) of the Universe is defined by the full 3-dimensional matter density field $\delta(\mathbf{x}, t)$. Although it is difficult to determine $\delta(\mathbf{x}, t)$ directly, we extract information about it indirectly in two general ways: (i) light from distant sources (including the cosmic microwave background) is impacted by over and underdense regions; and (ii) gravitationally bound objects such as galaxies and clusters often *trace* the matter density. Examples of the first class of information include the late-time integrated Sachs-Wolfe (ISW) effect caused by decaying gravitational potentials in the dark-energy era and the deflection of photons due to gravitational lensing. The second class includes galaxy clustering and cluster counts. One important objective of modern cosmology is to develop statistical methods to combine this information in the most efficient and consistent manner, in order to reconstruct $\delta(\mathbf{x}, t)$ and constrain models of its origin and evolution.

In the past decades, independent experiments have made extraordinary advances in charting these individual tracers. For example, on the cosmic microwave background (CMB) front, several generations of anisotropy and polarization measurements have led to recent results; the Planck Collaboration has mapped the temperature and polarization anisotropy of the early Universe and used its lensing statistics to study the integrated gravitational potential along the line of sight [1,2]. The Atacama Cosmology Telescope (ACT) and the South Pole Telescope (SPT) have made similar achievements with smaller footprints but higher resolutions [3–6]. Stage-III wide-field photometric surveys such as the Dark Energy Survey (DES), the KioDegree

Survey (KiDS) and the Hyper Suprime-Cam (HSC) have observed millions of galaxies on a significant fraction of the sky and used galaxy positions and shape statistics to probe the low-redshift matter distributions [7–9]. The recipe for analyzing most of this data involves first converting the data into 2-dimensional maps (e.g., for CMB surveys) and catalogs (e.g., for galaxy surveys), computing the correlation functions (or the power spectra) of these fields, and then comparing these observed correlation statistics to a cosmological model in a Bayesian likelihood analysis to yield cosmological parameter constraints. In almost all of these cases, the fiducial cosmological model, Λ CDM, fits the data well.

In addition to these results from single probes, there has been an increased effort to maximize information by combining probes. An example of this is the recent DES result combining its data of galaxy positions and galaxy shapes with the projected gravitational potential measured by SPT and Planck [10]. In this example, roughly the same recipe is followed: DES made maps of the galaxy density in five tomographic bins and the shear in four bins; these were combined with maps of the projected gravitational potential from SPT and Planck. Given these three sets of maps, there are six sets of two-point functions (galaxy clustering, galaxy-galaxy lensing, cosmic shear, cosmic shear \times CMB lensing, galaxy density \times CMB lensing, and the CMB lensing autocorrelation function). This set of six two-point functions forms the data vector, which is then used to constrain parameters. The main goal of this effort is to extract from all this low-redshift (much lower than the decoupling of the CMB) data a measurement of the amount of clustering at late times. This is often quantified with S_8 ,

which the DES + SPT analysis determined to be $S_8 = 0.792 \pm 0.012$, lower than the Planck measurement, $S_8 = 0.832 \pm 0.013$. The discrepancy does not meet strict statistical standards but it has spawned much interest and it is reminiscent of the Hubble tension that is driven by different measurements of the zeroth order expansion rate of the Universe.

Taking stock, the fiducial cosmological model fits most of the data, but there are alluring hints that it is flawed, and one of the most intriguing ways of stress testing the model is to measure how the clustering of matter evolves over the course of time. To date, this has been done predominantly by: (i) map-making; (ii) compression to two-point functions; and (iii) parameter constraints.

Research into field-level analysis offers an opportunity to change the way that we extract data from surveys, in the process offering an alluring opportunity for a powerful suite of tests of Λ CDM. The basic idea of field-level analysis is to combine all three steps above into one. Early examples of this idea [11–16] focused on the CMB. In that example, the time-ordered data can be converted into a map at the same time that the power spectrum is determined. The parameters to fit for the data, therefore, are the values of the temperature in all the pixels in the map plus a handful of cosmological parameters that determine the power spectrum. Eriksen *et al.* [14] extended the idea to allow for multiple maps to be constructed; e.g., maps of foregrounds in addition to the CMB. This basic technology has been incorporated into the most recent results from Planck [17]. Groups are now applying the technology to galaxy surveys [18–23].

One way to understand the advantage of the field-level approach is to return to the DES + SPT analysis; first CMB lensing maps were made using the traditional quadratic estimator [24] and then they were used to construct two-point functions. However, the data in DES itself could in principle help improve the fidelity of the CMB lensing maps; after all, the deflection of the CMB photons is due (at least in part) to the very structure that DES measures. Combining this information would clearly create a better CMB lensing map. Using that improved map with DES maps though would be a form of double counting, so it makes sense to do everything at once; create all the maps and estimate all the power spectrum simultaneously. In the particular example of CMB lensing, the problem is not trivial but Millea *et al.* [25–27] have made significant progress simultaneously measuring the lensing field, the primordial CMB, and several parameters that determine the relevant power spectra.

Here we use simulated data sets on large scales to: (i) develop the machinery that can handle real data; (ii) explore some of the basics of field-level analyses; and (iii) provide an example of how the field-level analyses can be used to stress test Λ CDM. Our example is related to the work in Eriksen *et al.* [14], except that we attempt to separate the late-time integrated Sachs-Wolfe signal from

the primordial CMB anisotropies. Hang *et al.* [28] constrained the ISW and lensing amplitudes using the two-point correlation between the DESI Legacy Survey and Planck temperature and lensing maps, where as we are interested in the full posterior distribution of both the parameters, two-point functions, and the maps.

We begin in Sec. II by explaining some of the details and then in Sec. III, we analyze simulated CMB data assuming that it consists only of noise and CMB anisotropies. We recover some of the known problems of the maximum posterior solution (the Wiener filter) and show that these can be mitigated by instead using samples of the full posterior. Then, in Sec. IV, we introduce the ISW component and try to separate that from the primordial anisotropies. The degeneracies make this problematic at the map level, but the sampler produces an unbiased power spectrum for each. This is crucial, as the cosmological parameters themselves are embedded in the spectrum so if the spectrum is unbiased, then the parameters will be as well. Specifically, we introduce two free amplitudes of each component that multiply the fiducial spectra and show that the field-level analysis that simultaneously solves for the map values and the parameters produces unbiased estimates of the parameters.

The ensuing constraints on the amplitude of the ISW spectrum are not very restrictive, so in Sec. V, we explore the possibility of adding in other tracers, the galaxy density in several tomographic bins. This adds to the number of free parameters in the field-level analysis but we show that it produces a higher-fidelity ISW map and a fairly tight constraint on the amplitude of the ISW spectrum. This leads to the prospect of stress testing Λ CDM by introducing amplitudes in front of all spectra (CMB lensing, galaxy density, cosmic shear) in addition to the standard cosmological parameters; a measurement in which any one of these amplitudes is determined to deviate from unity will disprove Λ CDM by demonstrating that structure does not grow in time as predicted by the model.

In short, our goal in this paper is to explain (to some, much of this will not be new) what to expect when carrying out a field-level analysis; demonstrate how well it does on simulated data with increasing numbers of components and probes, and point the way to a simple but powerful way to stress test the fiducial cosmological model. We share our conclusions and thoughts about the next steps in Sec. VI.

II. THEORETICAL FRAMEWORK

A. Field-level multiprobe analysis

The general problem of a field-level multiprobe inference is summarized in Fig. 1. The data is an observation, or a set of observations, on the sky. For concreteness, we will focus on the synergy between CMB experiments and photometric galaxy surveys, but the argument generalizes to any combinations of probes.

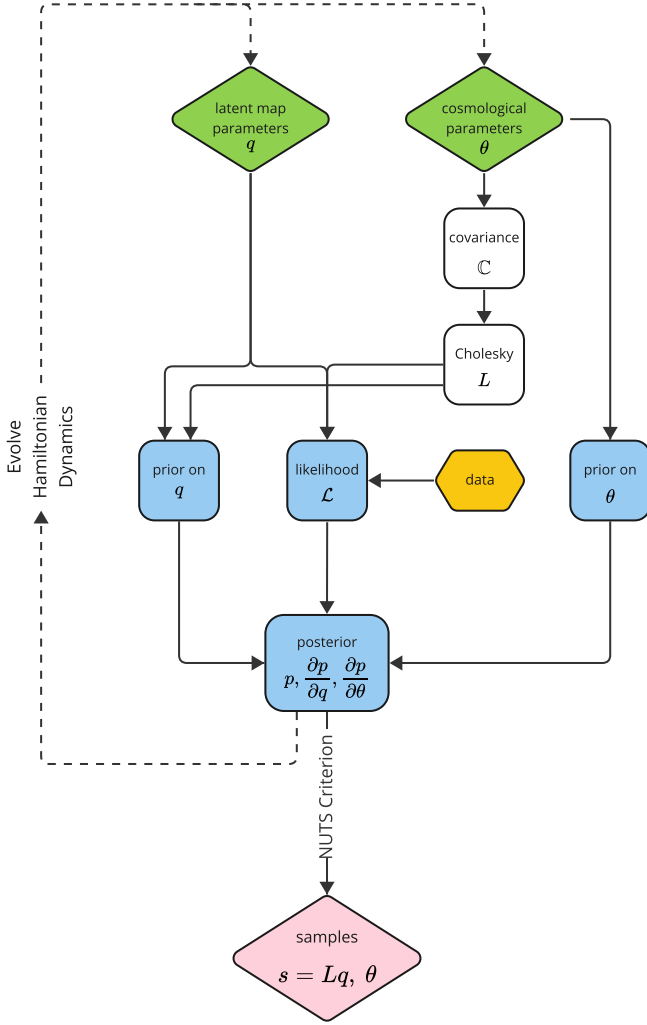


FIG. 1. The flow chart of a general field-level multiprobe analysis that accumulates samples of both the cosmological parameters and the tracer maps (values of each signal in each pixel) as discussed in Sec. II A and Sec. II C. We start from the prior distribution of the cosmological and latent map parameters. We then use the realized cosmological parameters to construct the covariance of the tracers, which in turn transforms the latent map parameters into physical tracer maps. The covariance, tracer maps, and observed data are then combined into the likelihood and—after multiplying by the priors—the posterior functions. If the above calculations are all programmatically differentiable, we can calculate the derivatives of the posterior function easily, and use HMC No-U-Turn Sampler (HMC-NUTS) to efficiently sample from the very high-dimensional posterior space. In this diagram, diamonds denote sampled parameters, squares denote model-relevant functions, and the hexagon is the (fixed) observed data vector. The pink diamond represent samples that represents the posterior space.

The data is assumed to consist of a set of signals and noise,

$$d = \sum_{\alpha} s^{\alpha} + n. \quad (1)$$

Our model assumes that the signals s^{α} in the data are drawn from a Gaussian distribution with mean zero and covariance matrix, $\mathbb{C}^{\alpha\beta}(\theta)$, where θ represents cosmological and nuisance parameters. The noise is also drawn from a Gaussian distribution with mean zero and known covariance matrix \mathbb{C}^n .

The likelihood for obtaining the data given the cosmological parameters and the signals is

$$-2 \ln \mathcal{L} = \left[d - \sum_{\alpha} s^{\alpha} \right] [\mathbb{C}^n]^{-1} \left[d - \sum_{\beta} s^{\beta} \right] + \dots, \quad (2)$$

where the additional terms are irrelevant, and the products on the right involve all pixels. That is, in the case of a single survey with N_{pix} pixels, d is a set of the values in all the pixels, and \mathbb{C}^n is a $N_{\text{pix}} \times N_{\text{pix}}$ matrix. If only one signal contributes, then s also has N_{pix} values; if more signals are assumed, then the total number of parameters in all the s^{α} will be $N_{\text{pix}} \times N_{\text{signal}}$. When data from multiple surveys are used, d will be a concatenated version of all the individual data sets and different signals can contribute to different data sets.

Using Bayes theorem, we can invoke the prior on all the signals and the parameters. Since we are confining our analysis to large scales throughout, the prior on all signals is Gaussian, and the posterior is

$$-2 \ln p = -2 \ln \mathcal{L} + \sum_{\alpha\beta} s^{\alpha} (\mathbb{C}^{-1}(\theta))^{\alpha\beta} s^{\beta} + \ln \det \mathbb{C}(\theta) - 2 \ln \text{prior}(\theta), \quad (3)$$

where irrelevant terms have been dropped. The parameters in this posterior are θ (which determines \mathbb{C}) and the map(s) s^{α} . For example, in the case of a single survey, if there is one signal contributing and there are five cosmological parameters, then the number of parameters we use to fit the N_{pix} data points is $N_{\text{pix}} + 5$. There are often cases where there are two or more signals contributing. For example, below we model the CMB as consisting of the signal from the last scattering surface plus the contribution from the late-time ISW effect. In that case, there will be $2N_{\text{pix}} + 5$ free parameters.

As described in Sec. II C, we will draw samples from this posterior. The accumulated samples of both the maps and the cosmological and nuisance parameters are fully consistent in the Bayesian sense. More precisely, the distribution of the values of map pixels s^{α} will provide a set of posterior samples of the signals, and the distribution of the parameters will constrain the relevant models of interest. These distributions will be consistent with one another, so that for example in a sample with a large $\mathbb{C}^{\alpha\alpha}$, the signal s^{α} everywhere is likely to have a larger dispersion.

B. Pixels

Above we glossed over the details of the map. Here, we review the basics of pixels in terms of the coefficients of spherical harmonics and explain why we choose to work with this basis.

Consider a map on the curved sky $s(\mathbf{n})$, where \mathbf{n} is a 3-dimensional unit vector. Analogous to Fourier transformations in Euclidean spaces, we can study this field in the frequency (or harmonics) space via forward and inverse spherical harmonics transform (SHT),

$$s(\mathbf{n}) = \sum_{lm} s_{lm} Y_{lm}(\mathbf{n}) \quad (4)$$

$$s_{lm} = \int \frac{d\Omega}{4\pi} s(\mathbf{n}) Y_{lm}^*(\mathbf{n}), \quad (5)$$

where $Y_{lm}(\mathbf{n})$'s are the set of orthonormal spherical harmonics. We adopt the HEALPix pixelization strategy (where the angular resolution is specified by a single parameter NSIDE), and use the discretized SHT as implemented by the HEALPY library [29,30]. As usual in cosmological analyses, we drop the monopole and dipole modes ($l = 0, 1$).

In general, if the field s is statistically homogeneous and isotropic, it is more advantageous to study s 's correlation structure in harmonic space. In real space, the correlation function between two line-of-sight directions is given by

$$w_{\mathbf{n},\mathbf{n}'} = w(|\mathbf{n} - \mathbf{n}'|) = \langle s(\mathbf{n})s(\mathbf{n}') \rangle, \quad (6)$$

where we see that the correlation function has dense off-diagonals. For a discretized map with NSIDE resolution, the size of w scales as NSIDE^4 , which quickly becomes impossible to handle (for example, an NSIDE = 256 map has an angular resolution of $27''$ and 8×10^5 pixels; the full pixel-pixel covariance matrix totals 5 terabytes).

However, s 's power spectrum (s 's correlation function in harmonic space), \mathbb{C} , defined by

$$\langle s_{lm}^\alpha s_{l'm'}^{*\beta} \rangle = \delta_{ll'} \delta_{mm'} \mathbb{C}_l^{\alpha\beta} \quad (7)$$

is diagonal in this basis and depends only on the multiple moment l and the different sets of signals assumed. Therefore, the amount of memory needed to manipulate \mathbb{C} is linear in NSIDE. One important caveat to this simplicity is that the field must be homogeneous and isotropic, and these assumptions fail in the presence of instrumental noise patterns, partial sky coverage, and masking.

C. Methodology

Here we present the details of our implementation of Fig. 1. The fundamental idea behind all MC sampling techniques is to start from the current sample and then find the next point in the parameter space and generate a probabilistic proposal to make it a sample (both operations may involve repeated evaluations of the posterior density). How to find the next point and what proposal to make are algorithm specific; however, they in general satisfy the principle of detailed balance such that, in the limit of large sample size, the samples approximate the posterior distribution.

The efficiency of MC sampling rests on the suitability of the MC algorithm for the specific inference context and the effective computation of the posterior distribution.

For the first point, since we are inferring both the map pixels and the cosmological parameters, the dimensionality of the posterior space will be quite large. For example, for our final analysis in Sec. V which includes eight tracer maps at NSIDE = 32, the total dimensionality of the posterior space is 73704. This is too large for traditional Monte Carlo techniques (such as Metropolis-Hastings) to operate efficiently. Intuitively, this is because as dimensionality increases, the ratio between the neighboring volume pointing towards and away from a particular point in the parameter space (e.g., the mode of the distribution) decays exponentially. Thus, the random walk Metropolis algorithm becomes overwhelmingly likely to propose samples outside the typical set, where the target density and hence the acceptance probability vanishes [31].

Hamiltonian Monte Carlo (HMC) solves this efficiency problem in high-dimensional spaces [31–33]. In the HMC framework, we augment the parameter space with a conjugate momentum space and use the gradient of the log posterior surface to guide us to sample only near the bulk of the probabilistic mass [31,32]. In order to avoid traditional HMC's sensitivity to hyper-parameters such as the integration steps, we further employ the No-U-Turn Sampler (NUTS) variation of the HMC, first proposed by Hoffman and Gelman [33].

This leads us to the second point on computational efficiency. HMC samplers require repeated evaluations of the posterior function and its gradient. Since we want to develop a multiprobe field-level framework that easily extends to different observables and cosmological models, we do not want to hard code the derivatives in advance. Instead, we choose to make the framework pragmatically differentiable through the JAX autodifferentiation library in Python, which interfaces smoothly with the numpyro implementation of the NUTS [34,35].

Turning to the specific problem of the posterior function computation. We start at the top of Fig. 1 and break this calculation into several parts:

- (i) We start with a proposal for the cosmological parameters θ and the latent map parameters. The latent map parameters are a set of uncorrelated standard Gaussian variables q which we will later transform into the signal maps.
 - (ii) Calculate \mathbb{C} given the cosmological parameters
 - (iii) Transform the latent map parameters using the Cholesky decomposition ($\mathbb{C} = LL'$) of the covariance matrix, $s = Lr$. [The prior of the maps becomes $s^2/(2\mathbb{C}) = q^2/(2\mathbb{I})$].
 - (iv) Combine the maps s with the data to calculate the likelihood (forward modeling)
 - (v) Use the likelihood and the prior to calculate the posterior and its derivative with respect to the parameters
 - (vi) If the NUTS criterion is satisfied, accept this as a valid sample
 - (vii) Use the leapfrog method to generate another sample
- In practice, through the JAX autodifferentiation library, this framework provides information on both the posterior and its gradient.

The treatment of the covariance function deserves more discussion. In this paper, we will keep the shape of the spectra fixed and allow for free amplitudes $(A^\alpha)^2$. We assume the amplitudes have fiducial values equal to one and have a uniform prior distribution. We then construct the full covariance matrix $\mathbb{C}^{\alpha\beta}$, which consists of the auto and cross-spectra of each signal. In principle, since \mathbb{C} encodes the covariance between all the pixels for all the tracer maps, its dimensionality is very high. For a single HEALPix map at the resolution of NSIDE (or a limiting resolution of $l_{\max} = 3\text{NSIDE} - 1$), there are

$$\sum_{l=2}^{l_{\max}} (2l+1) = l_{\max}^2 + 2l_{\max} - 3 \quad (8)$$

degrees of freedom ignoring monopole and dipole modes. Again, for our final analysis in Sec. V, which includes eight tracers, the size of \mathbb{C} is on the order of $8^2 \times l_{\max}^4$. The efficient computation of this covariance matrix is one of the limiting factors in the feasibility of field-level analysis. However, in the limit of full sky and when all the fields are homogeneous and isotropic, the sub-block of \mathbb{C} for each tracer is diagonal in the a_{lm} basis. Thus, we can bring \mathbb{C} into block diagonal forms, with $l_{\max} - 2$ unique \mathbb{C}_l sub-blocks on the diagonal. Each \mathbb{C}_l sub-block has size 8×8 , describing the correlation between the 8 tracers at mode l . Looking ahead, we will be considering the primordial CMB signal (modulated by $A^{\mathcal{P}}$); the late-time ISW effect (modulated by $A^{\mathcal{I}}$) and the galaxy density in six tomographic bins (modulated by b_i 's). The first of these is uncorrelated with the rest, so the ensuing 8×8 sub-block matrix will be

$$\mathbb{C}_l(b_1, \dots, b_6, A^{\mathcal{I}}, A^{\mathcal{P}}) = \begin{pmatrix} b_1^2 \mathbb{C}_l^{\mathcal{G}_1, \mathcal{G}_1} & b_1 b_2 \mathbb{C}_l^{\mathcal{G}_1, \mathcal{G}_2} & \dots & b_1 A^{\mathcal{I}} \mathbb{C}_l^{\mathcal{G}_1, \mathcal{I}} & 0 \\ b_1 b_2 \mathbb{C}_l^{\mathcal{G}_2, \mathcal{G}_1} & b_2^2 \mathbb{C}_l^{\mathcal{G}_2, \mathcal{G}_2} & \dots & b_2 A^{\mathcal{I}} \mathbb{C}_l^{\mathcal{G}_2, \mathcal{I}} & 0 \\ \dots & \dots & \dots & \dots & \dots \\ b_1 b_6 \mathbb{C}_l^{\mathcal{G}_6, \mathcal{G}_1} & b_2 b_6 \mathbb{C}_l^{\mathcal{G}_6, \mathcal{G}_2} & \dots & b_6 A^{\mathcal{I}} \mathbb{C}_l^{\mathcal{G}_6, \mathcal{I}} & 0 \\ b_1 A^{\mathcal{I}} \mathbb{C}_l^{\mathcal{I}, \mathcal{G}_1} & b_6 A^{\mathcal{I}} \mathbb{C}_l^{\mathcal{I}, \mathcal{G}_2} & \dots & (A^{\mathcal{I}})^2 \mathbb{C}_l^{\mathcal{I}, \mathcal{I}} & 0 \\ 0 & 0 & \dots & 0 & (A^{\mathcal{P}})^2 \mathbb{C}_l^{\mathcal{P}, \mathcal{P}} \end{pmatrix}. \quad (9)$$

With this computationally efficient representation of the covariance matrix, we can transform the latent map variables into the proper tracer maps through either the sub-blocks' Cholesky representations or their eigendecomposition. We experimented with both, and found the former to be an order of magnitude faster (see also Loureiro *et al.* [19]). In summary, the algorithm is very fast. For the largest model we considered in Sec. V, the analysis was done on an Apple M1 chip running overnight.

III. RECONSTRUCTION OF THE PRIMORDIAL CMB MAP

We start with the simplest possible example, a single simulated CMB all-sky map. Although simple, this model demonstrates the key behaviors of two ways of using the posterior; identifying the free parameters by finding the point at which the posterior is maximum (hereafter, maximum *a posteriori* or MAP) and generating samples of the posterior (hereafter sampling). We compare the potential biases of both methods and discuss the implication for the field, two-point, and cosmological parameter constraints. In some ways, the idea of asking whether an estimator is biased is introducing frequentist ideas into a Bayesian discussion. Nonetheless, we think that understanding these biases is an important step towards the ultimate goal of extracting the correct cosmological parameters from the data. The intuition we find here will serve us well in the subsequent more complex cases.

The primordial CMB temperature fluctuations $s^{\mathcal{P}}$ originate from the time of recombination ($z_* \approx 1100$), when the photons decoupled from the photon-electron-proton fluid as the universe cooled below a few percent of the ionization energy of hydrogen. We assume (and all current data is consistent with this assumption, with the tightest constraints coming from the Planck Collaboration *et al.* [36]) that the resulting temperature variation is a homogeneous and isotropic random Gaussian field which is fully characterized by the power spectrum $\mathbb{C}_l^{\mathcal{P}}$. Throughout this paper, we assume a fiducial cosmology of $H_0 = 100h = 67.5$ km/Mpc/sec, $\Omega_b h^2 = 0.0219$, $\Omega_c h^2 = 0.1139$, $A_s = 2 \times 10^{-9}$ and $n_s = 0.965$.

A. Biases of the optimal estimator

1. Fixed cosmological parameters

The observed temperature data $d^T(\mathbf{n})$ is the superposition of the primordial field $s^P(\mathbf{n})$ and noise $n^T(\mathbf{n})$.

$$d^T(\mathbf{n}) = s^P(\mathbf{n}) + n^T(\mathbf{n}) \quad (10)$$

We call this the CMB model. We ask; Given d^T and perfect knowledge of s^P 's and n^T 's theoretical power spectra (C^P and $C^{n,T}$ respectively), how well can we reconstruct the primordial field? Additionally, how accurate is the power spectrum of the reconstructed field?

In the Bayesian framework, the posterior probability in Eq. (3) reduces to

$$-2 \ln p(s^P | d^T) \propto \sum_{lm} \left(\frac{|d_{lm}^T - s_{lm}^P|^2}{C_l^{n,T}} + \frac{|s_{lm}^P|^2}{C_l^P} \right), \quad (11)$$

where we drop the determinant terms since the CMB and noise spectra are assumed known and fixed. The MAP solution for s^P is then given by the Wiener filter

$$\hat{s}_{lm}^{P,\text{MAP}} = \frac{C_l^P}{C_l^P + C_l^{n,T}} d_{lm}^T. \quad (12)$$

The mean power spectrum of the MAP estimator

$$\langle \hat{C}_l^{P,\text{MAP}} \rangle \equiv \left\langle \frac{1}{2l+1} \sum_m |\hat{s}_{lm}^{P,\text{MAP}}|^2 \right\rangle = \frac{C_l^P}{C_l^P + C_l^{n,T}} C_l^P \quad (13)$$

is known to be biased [37,38], an effect more prominent in the low signal-to-noise ratio (SNR) regime.

To implement, we simulate s^P from the fiducial cosmology power spectrum on a HEALPix grid of $\text{NSIDE} = 64$ and then inject isotropic white noise n^T with a relatively high variance of $\text{Var}(n^T) = 1000 \mu\text{K}^2$. Even though this exceeds noise in Planck by several orders of magnitude, we use this value to demonstrate the difficulties of extracting the signal in the presence of appreciable noise. The power spectra of the truth map (black) and the recovered MAP map (green) are shown in Fig. 2. As the amplitude of the noise spectrum (purple) rises on small scales, the MAP spectrum is increasingly suppressed. On the field level, this means that $\hat{s}_{lm}^{P,\text{MAP}}$ is damped for over and underdensities on scales that have small SNR. The estimator does not have an additive bias but does have a multiplicative bias, i.e., $\langle \hat{s}_{lm}^{P,\text{MAP}} / s_{lm}^P \rangle = C_l^P / (C_l^P + C_l^{n,T})$ for both the real and the imaginary components. This is shown in the top panel of Fig. 3.

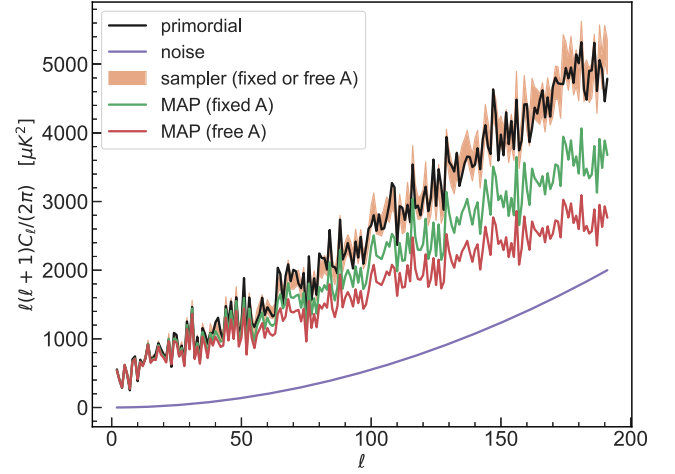


FIG. 2. A numerical experiment of CMB reconstruction that demonstrates the difference between the MAP estimators and the sampling-based estimators. We generated the truth map from a fiducial C^P and added isotropic noise. The truth and noise spectra are shown in black and purple. We test two reconstruction models, the first has a fixed C^P [Eq. (11)], and the second has a free amplitude A^2 modulating the fiducial C^P . Their MAP solutions are given by Eqs. (13) and (15), and are shown in green and red respectively. We confirm the analytical solutions match the results from direct numerical optimization. Both two-point and MAP estimators are biased lower than truth and the free amplitude model has a greater bias. We then sampled s^P directly from the posterior space. The spectra of the mean maps for the fixed and free amplitude cases are biased (both overlap the green curve). However, in both cases, the distributions (the orange shaded region represent 1σ credible interval) of the spectra scatter around the truth and are unbiased.

2. Varying cosmological parameters

In real cosmological analyses, we are also interested in cosmological parameters (such as the primordial amplitude, spectral index, etc.) that modify the shape and amplitude of the power spectrum. We want to know how the field, two-point, and parameter MAP estimators behave when the spectrum is allowed to change.

For example, consider modulating the fiducial power spectrum C^P with a scale-invariant amplitude A^2 , where A has a flat prior. The new MAP solutions are given by

$$\hat{s}_{lm}^{P,\text{MAP}} = \frac{(\hat{A}^{\text{MAP}})^2 C_l^P}{(\hat{A}^{\text{MAP}})^2 C_l^P + C_l^{n,T}} d_{lm}^T \quad (14)$$

$$\hat{C}_l^{P,\text{MAP}} = \frac{(\hat{A}^{\text{MAP}})^2 C_l^P}{(\hat{A}^{\text{MAP}})^2 C_l^P + C_l^{n,T}} (\hat{A}^{\text{MAP}})^2 C_l^P \quad (15)$$

and \hat{A}^{MAP} satisfies




FIG. 3. The scatter plot of the ratio between the real components of the MAP a_{lm} and that of the truth a_{lm} as a function of l . The cases for fixed and varying amplitude parameters are shown in the upper and lower panels respectively. In each case, the mean of the ratios (black curve) matches the Wiener filter expectations (with the MAP amplitude in the case of varying amplitude). The color map represents point cloud density.

$$\sum_l (2l+1) \left\{ 1 - \frac{(\hat{A}^{\text{MAP}})^2 \mathbb{D}_l \mathbb{C}_l^{\mathcal{P}}}{((\hat{A}^{\text{MAP}})^2 \mathbb{C}_l^{\mathcal{P}} + \mathbb{C}_l^{n,T})^2} \right\} = 0, \quad (16)$$

where \mathbb{D}_l is the data power spectrum.

When noise is present, $\hat{A}^{\text{MAP}} < 1$, in this case equal to 0.81 (see Fig. 4). Thus, by comparing Eqs. (13) and (15), we see that the new MAP is biased even lower than the truth. We can apply this model to the same reconstruction experiment as before. The result for the new power spectrum estimator is shown in red in Fig. 2. The multiplicative bias in the field-level estimator is shown in the bottom panel of Fig. 3.

For power spectra with complicated parameter dependence, we often lack analytical optimal solutions. However, qualitatively speaking, if an increase in the parameter increases the amplitude of the spectrum as in this case, then the parameter will be underestimated by optimal inference, and vice versa.

B. Sampling the CMB field

1. Fixed cosmological parameters

Now we seek an unbiased estimator for $\mathbb{C}^{\mathcal{P}}$ that also has a convenient notion of uncertainty. Let us again first fix $A = 1$ and draw a sample of maps $\{s_{lm,i}^{\mathcal{P}}\}_{i=1,\dots,N}$ directly from the posterior distribution [Eq. (11)]. Using this set of maps, we can construct an associated set of power-spectra samples

$$\{\mathbb{C}_{l,i}^{\mathcal{P}}\} = \left\{ \sum_m \frac{1}{2l+1} |s_{lm,i}^{\mathcal{P}}|^2 \right\}. \quad (17)$$

Let us call the ensemble average of $\{s_{lm,i}^{\mathcal{P}}\}$ and $\{\mathbb{C}_{l,i}^{\mathcal{P}}\}$ as $\overline{s^{\mathcal{P}}}$ and $\overline{\mathbb{C}^{\mathcal{P}}}$, respectively. It is crucially important that the power spectrum of $\overline{s^{\mathcal{P}}}$ is *different* from $\overline{\mathbb{C}^{\mathcal{P}}}$.

We claim, in the limit of sufficient sample size N ,

- (1) $\overline{s^{\mathcal{P}}}$ and its spectrum are exactly the field and two-point MAP estimators [Eqs. (12) and (13)], and they have the Wiener filter multiplicative bias.
- (2) On the two-point level, the samples $\{\mathbb{C}_{l,i}^{\mathcal{P}}\}$ give a proper Bayesian credible interval centered around the truth.
- (3) Further, $\overline{\mathbb{C}^{\mathcal{P}}}$ (and more generally, the mean of any n -point power spectrum samples) is an *unbiased* estimator in the frequentist sense (when we have multiple data realizations).

We prove these claims in Appendix B. However, intuitively, how can the power spectrum of the mean map be biased while the sampled power spectrum be unbiased? One way to understand this is to think of each sampled map as $s^{\mathcal{P}} = \hat{s}^{\mathcal{P},\text{MAP}} + s'$. When we compute the power spectrum $\langle |s^{\mathcal{P}}|^2 \rangle$, the $\langle |s'|^2 \rangle$ term exactly compensates for the deficiency of the MAP spectrum. Alternatively, the $s^{\mathcal{P}}$'s are normally distributed, and for any Gaussian distribution the mean is equal to the maximum, so $\overline{s^{\mathcal{P}}}$ is the MAP solution. However, the power spectrum is not normally distributed; its expected value is an unbiased estimator of truth, and not the biased MAP solution.

We continue with the numerical experiment above. This time, we construct an HMC NUT sampler following the prescription of Sec. II C, using Eq. (11) as our posterior distribution. After the chain equilibrates, we draw 3000 $s^{\mathcal{P}}$'s from the posterior space. We confirm that the power spectrum of the mean map exactly follows the MAP solution for the case of fixed parameters (the green curve in Fig. 2). We further show the distribution of the sampled power spectra $\{\mathbb{C}_{l,i}^{\mathcal{P}}\}$ in Fig. 2 in shaded orange. Indeed, the distribution of the spectra covers the truth power spectrum within uncertainty, and $\overline{\mathbb{C}^{\mathcal{P}}}$ is unbiased.

We note that similar phenomena have been observed in previous studies. For example, in Fig. 8 of [26], the authors find that the distribution of the sampled CMB spectra scatter around truth while the spectrum of the mean map is biased lower at small scales.

The above observations have the following implications. One must debias the sampled maps before using them for cosmological analysis, similar to how we currently correct for MAP maps (e.g., with analytical or Monte Carlo-based corrections). However, if we are only performing analysis on the two-point level, the samples are *unbiased* and their distribution constitutes a convenient measure of uncertainty.

In short, by considering the sampled power spectra, we recapture an unbiased estimator of $\mathbb{C}^{\mathcal{P}}$.

2. Varying cosmological parameters

The exact two-point statistics recovery motivates us to ask whether the sampled cosmological parameters that modify the power spectrum are also unbiased. To answer this question, we use the HMC NUT sampler from the previous section with an additional A^2 modulating the fiducial spectrum. We assume A has a flat and wide prior on $[0.2, 10]$ and collect 3000 samples after appropriate burn-in.

We find that A is unbiased with its marginal distribution shown in Fig. 4. We can also estimate the variance of its distribution, which is predicted by the inverse of the Fisher information

$$\mathcal{F}^{-1} = \sum_l \frac{2(2l+1)}{1 + (\mathbb{C}^{n,T}/\mathbb{C}^{\mathcal{P}})^2}, \quad (18)$$

shown as shaded orange in the same figure.

Since the expectation of A is unbiased, it follows that the posterior samples $\{\mathbb{C}_{l,i}^{\mathcal{P}}\}$ again scatter around truth unbiased. In fact, their distribution overlaps that of the fixed amplitude model almost exactly as shown in Fig. 2 in shaded orange. Further, the power spectrum of the mean map is also not the MAP anymore; it is the MAP solution of the model with fixed cosmological parameter [Eq. (13)], as if A is fixed to 1.




FIG. 4. The posterior distribution of the amplitude A (blue curve) compared to truth (black dashed line). The shaded orange area represents the Fisher forecast of the 1σ uncertainty centered on truth, and the shaded blue area represents the 1σ uncertainty reported by the sampler. The red line shows the MAP value for A , which is quite far from truth.

The important takeaway is the following. The MAP map and the mean sampled map are biased both on the field level and on the two-point level. However, the distribution of the sampled power spectra (and cosmological parameters that modulate them) is unbiased.

IV. JOINT RECONSTRUCTION OF THE PRIMORDIAL CMB AND THE ISW EFFECT

Now, we expand on the CMB model and consider extracting the primordial and the ISW contributions from a single noisy temperature measurement. As we shall see, the MAP estimators for both signals are again biased on the field and the two-point level. The sampled fields have a multiplicative bias but their two-point statistics are unbiased. The new challenge in this case study is the field-level degeneracy between the primordial and the ISW maps, which motivates the multiprobe approach presented in the next section [14]. We discuss the key properties of this degeneracy, which we expect to be quite general when one separates a low SNR map from a measurement based on a likelihood approach.

A. The ISW effect

In the late universe, the primordial CMB fluctuations are modified by the ISW effect on very large scales [39–43]. A photon is blue-shifted when descending a gravitational potential well and red-shifted when it escapes. When the universe began its accelerated expansion (in the dark energy-dominated era), the large-scale potential wells decayed. As a result, a photon will leave the decaying well (barrier) with more (less) energy than it enters. The observed ISW temperature modification is thus the integrated effect of the decaying potential well along the line of sight, and its 2-dimensional field is given by [44,45]

$$A^{\mathcal{I}}(\mathbf{n}) = \int_0^\infty \frac{\partial \Phi(\mathbf{x}[\chi, \mathbf{n}], t(\chi))}{\partial t} \frac{2e^{-\tau(\chi)}}{1+z(\chi)} d\chi. \quad (19)$$

Here, $\tau(\chi)$ is the optical depth out to distance χ and Φ is the 3-dimensional gravitational potential field, which ultimately depends on the matter overdensities $\delta_m(\mathbf{x}[\chi, \mathbf{n}], t(\chi))$. By rewriting Φ in terms of δ_m , and moving to the Fourier space, we can rewrite Eq. (19) as

$$a_{lm}^{\mathcal{I}} = 4\pi i^l \int \frac{d^3k}{(2\pi)^3} I_l^{\mathcal{I}}(k) Y_{lm}^*(\mathbf{k}) \delta_m(\mathbf{k}, t_0), \quad (20)$$

where

$$I_l^{\mathcal{I}}(k) = \int d\chi D(\chi) W^{\mathcal{I}}(k, \chi) j_l(k\chi) \quad (21)$$

and the window function is given by

$$W^{\mathcal{I}}(k, \chi) = -\Theta(\chi_* - \chi) \frac{3\Omega_m H_0^2}{k^2} \frac{\partial \ln((1+z)D(z))}{\partial t}, \quad (22)$$

where $D(z)$ is the growth function normalized to unity at $z = 0$ and we approximate τ as zero through the epoch that the ISW is generated.

The forms of Eqs. (20)–(22) are not peculiar to the ISW effect—by modifying the window function $W(\chi, k)$ appropriately, the 2-dimensional observable of most tracers can be computed as a line-of-sight integral of δ_m . For general tracers, A and B , of this form, the covariance is

$$\mathbb{C}_l^{A,B} = \frac{2}{\pi} \int_0^\infty k^2 dk P(k) I_l^A(k) I_l^B(k). \quad (23)$$

One way of quantifying the correlation between different probes is to compute the scale-dependent correlation coefficients, defined as

$$\rho^{AB} = \frac{\mathbb{C}^{AB}}{\sqrt{\mathbb{C}^{AA}\mathbb{C}^{BB}}}. \quad (24)$$

We again assume both fields are statistically homogeneous and isotropic, so the covariance is diagonal in the a_{lm} basis. For example, the primordial CMB and the ISW effect are spatially independent ($\mathbb{C}^{\mathcal{I},\mathcal{P}} = 0$), and their autopower spectra are shown in Fig. 5 ($\mathbb{C}^{\mathcal{I},\mathcal{P}} = 0$). The ISW signal is primarily confined to the very large scales ($l < 10$) that enter the horizon during the dark energy-dominated era. The ISW signal is also subdominant to the primordial signal on all scales, making it particularly challenging to reconstruct. We explore the case of non-diagonal covariance in Sec. V in the context of multiprobe joint reconstruction.




FIG. 5. The power spectra of the primordial CMB and the ISW effect on large scales (low- l modes).

B. Separating the primordial and ISW signals

The observed temperature data $d^{\mathcal{I}}(\mathbf{n})$ is the sum of the primordial field $s^{\mathcal{P}}(\mathbf{n})$, the ISW field $s^{\mathcal{I}}(\mathbf{n})$, and noise $n^{\mathcal{I}}(\mathbf{n})$, so the posterior distribution

$$\begin{aligned} \ln p(s^{\mathcal{P}}, s^{\mathcal{I}}, A^{\mathcal{P}}, A^{\mathcal{I}} | d^{\mathcal{I}}) \\ \propto \sum_{lm} \frac{1}{2} (\log A^{\mathcal{P}2} \mathbb{C}_l^{\mathcal{P}} + \log A^{\mathcal{I}2} \mathbb{C}_l^{\mathcal{I}}) \\ + \sum_{lm} \left(\frac{|d_{lm}^{\mathcal{I}} - s_{lm}^{\mathcal{P}} - s_{lm}^{\mathcal{I}}|^2}{2\mathbb{C}_l^{n,\mathcal{I}}} + \frac{|s_{lm}^{\mathcal{P}}|^2}{2A^{\mathcal{P}2}\mathbb{C}_l^{\mathcal{P}}} + \frac{|s_{lm}^{\mathcal{I}}|^2}{2A^{\mathcal{I}2}\mathbb{C}_l^{\mathcal{I}}} \right). \end{aligned} \quad (25)$$

We will refer to this model as the CMB-ISW model.

One crucial difference between the CMB and the CMB-ISW model is that although both are constrained by the same amount of data $d^{\mathcal{I}}$, the dimensionality of the latter's posterior space ($s^{\mathcal{P}}$ and $s^{\mathcal{I}}$) is (ignoring the cosmological parameters) twice that of the former's (only $s^{\mathcal{P}}$). In other words, if we have N independent modes on the full sky, we are trying to constrain $2N$ parameters with N data points in the CMB-ISW model. Hence, we expect significant degeneracy in the inferred $s^{\mathcal{P}}$ and $s^{\mathcal{I}}$ maps, and we want to explore how this affects the MAP and the sampling-based field-level reconstructions. We emphasize that this problem will be quite common in any field-level analysis where we wish to separate the different physical components of a single observed field.

We will again tackle this problem in two ways, first by constructing the MAP estimators and then by sampling directly from the posterior distribution.

1. Fixed cosmological parameters

Let us first fix $A^{\mathcal{P}} = A^{\mathcal{I}} = 1$ and seek the field-level MAP solutions. When we are trying to find the MAP of $s^{\mathcal{P}}$, the effective noise is the sum of the instrumental noise and the ISW temperature fluctuation (and analogously for $s^{\mathcal{I}}$). Thus, invoking the Wiener filter [Eq. (12)],

$$\hat{s}_{lm}^{\mathcal{P},\text{MAP}} = \frac{\mathbb{C}_l^{\mathcal{P}}}{\mathbb{C}_l^{\mathcal{P}} + \mathbb{C}_l^{n,\mathcal{I}} + \mathbb{C}_l^{\mathcal{I}}} d_{lm}^{\mathcal{I}}, \quad (26)$$

$$\hat{s}_{lm}^{\mathcal{I},\text{MAP}} = \frac{\mathbb{C}_l^{\mathcal{I}}}{\mathbb{C}_l^{\mathcal{I}} + \mathbb{C}_l^{n,\mathcal{I}} + \mathbb{C}_l^{\mathcal{P}}} d_{lm}^{\mathcal{I}}, \quad (27)$$

and similarly for their power spectra.

As in Sec. III, we simulate this reconstruction method numerically by generating $s^{\mathcal{P}}$ and $s^{\mathcal{I}}$ with fiducial cosmology a HEALPix grid of $\text{NSIDE} = 32$ with $\text{Var}(n^{\mathcal{I}}) = 200 \mu\text{K}^2$. In the top panels of Fig. 6, the truth power spectra are shown in black, the MAP power spectra in green, and the effective noise in purple. The primordial MAP spectrum is biased low both on large scales (ISW contamination) and small scales (noise contamination). The ISW MAP spectrum

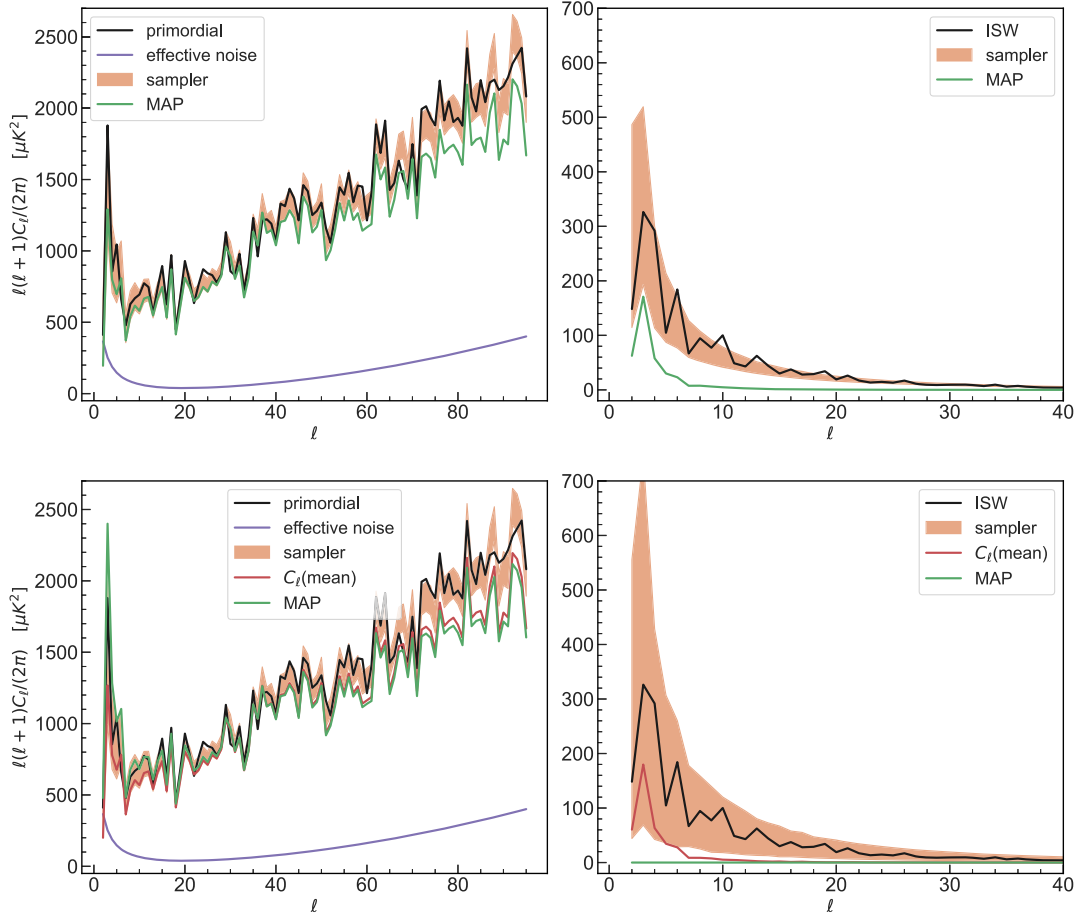


FIG. 6. The recovered primordial ($s^{\mathcal{P}}$) and ISW ($s^{\mathcal{I}}$) spectra using the various models in Sec. IV B. The simulation is drawn from the fiducial cosmology and a noise variance of $200 \mu\text{K}^2$. The top and bottom panels show the case of fixed and free amplitude respectively, and the left and right panels show the primordial and the ISW results, respectively. For each case, the truth spectrum is shown in black, the MAP result is shown in green, and the 1σ credible interval of the sampled spectra is shown in shaded orange. The effective noise of the primordial maps is $\mathcal{C}^{\mathcal{I}} + \mathcal{C}^{n,\mathcal{I}}$ shown in purple. The effective noise of the ISW map ($\mathcal{C}^{\mathcal{P}} + \mathcal{C}^{n,\mathcal{I}}$) is above the ISW spectra in the right panels on all scales and hence not shown in the plot. Notice that in the case of free amplitude, the MAP solution is different from the C_l of the sampled maps (red).

is significantly biased low on all scales due to the same Wiener filter suppression.

On the field level, we again expect (and indeed observe) no additive bias but a multiplicative bias on the $\langle \hat{s}_{lm}^{\text{MAP}}/s_{lm} \rangle$ proportional to the Wiener filter factor (for both the real and the imaginary components). The case for the ISW field is particularly egregious, as shown in the top panel of Fig. 7 (note that the y-axis does not even contain the unbiased case).

Now we turn to the reconstructed real space maps (rows 1–3 of Fig. 8) which shed more light onto the degeneracy between the reconstructed $s^{\mathcal{P}}$ and $s^{\mathcal{I}}$. Qualitatively, the primordial MAP map captures most features of the true signal, although the small-scale structures are suppressed due to Wiener filtering. However, the MAP estimator completely fails in the ISW reconstruction. In fact, the $\hat{s}_{lm}^{\mathcal{I},\text{MAP}}$ map looks like a low-pass filtered $s^{\mathcal{P}}$ map.

The physical explanation is that, when we observe a large-scale hot spot in the sky, it is impossible to confidently associate it with either the primordial CMB or the ISW effect since they both have high amplitudes at low l 's. However, when we optimize the posterior function with respect to $s^{\mathcal{I}}$, the algorithm neglects the $s^{\mathcal{P}}$ [Eq. (27)]. Thus, the algorithm inclines to increase the amplitude of the $s^{\mathcal{I}}$ map wherever we observe a large-scale hot spot in $d^{\mathcal{I}}$, even though it is most likely due to $s^{\mathcal{P}}$ since it has a greater power spectrum. As a result, the large-scale hot spots of the reconstructed $s^{\mathcal{I}}$ are heavily correlated with $s^{\mathcal{P}}$, even though they are spatially independent in theory. An analogous bias can be said for the reconstructed primordial map, i.e., the reconstructed primordial map is biased high where there is an ISW hot spot (although it is slightly more difficult to discern in the figure). In short, as we attempt to reconstruct two maps from a single observation using the MAP




FIG. 7. The scatter plot of the ratio between the MAP and the truth a_{im} 's real components. The upper panel shows the case with fixed amplitude and the bottom shows the case with free amplitude. In the case of the fixed amplitude, the mean of the ratios is again the Wiener filter factor (black). In the case of free amplitude, the MAP solution has a null spectrum. Thus the ratios scatter around 0 for all scales.

estimator, the degeneracy introduces significant bias on the field level that correlates with the two reconstructed maps.

Now, we apply the sampler, as defined in Sec. II C, with s^P , s^I as the free parameters. The two-point result is shown in the top panels of Fig. 6 in orange, where we observe that the spectra samples scatter around the truth unbiased, and the spectrum of the mean sampled map is equivalent to the MAP spectrum. The sampling result on the field level is shown in rows 1–3 of Fig. 8. Here we confirm that the mean sampled map is indeed the field-level MAP solution, and thus suffers from the same bias and degeneracy. Therefore, although the sampling approach solves the multiplicative bias on the two-point level, it is placing the right amount of power in the wrong place at the map level.

2. Varying cosmological parameters

We now vary the two amplitude parameters (A^P and A^I) in Eq. (25) and attempt to reconstruct them together with the map pixels. The MAP solutions for $\hat{s}^{P,\text{MAP}}$ and $\hat{s}^{I,\text{MAP}}$ are analogous to Eqs. (26) and (27), but with the signal power spectra multiplied by their MAP amplitudes squared, which satisfy

$$\sum_l (2l+1) \left\{ 1 - \frac{(A^{P,\text{MAP}})^2 \mathbb{D}_l C_l^P}{Q_l^2} \right\}, \quad (28)$$

$$\sum_l (2l+1) \left\{ 1 - \frac{(A^{I,\text{MAP}})^2 \mathbb{D}_l C_l^I}{Q_l^2} \right\}, \quad (29)$$

where

$$Q_l = (A^{P,\text{MAP}})^2 C_l^P + (A^{I,\text{MAP}})^2 C_l^I + C_l^{n,I}. \quad (30)$$

We apply this MAP estimator to the numerical experiment for the CMB-ISW model discussed above. The biased two-point results are shown in the bottom panels of Fig. 6 in green. In the case of the ISW reconstruction, the effective noise is so large that the slope of the posterior distribution with respect to the amplitude [Eq. (29)] never achieves 0. This results in A^I and hence the MAP spectrum being set to 0, which we confirm using direct numerical optimization. This effect is also shown in the bottom panel of Fig. 7 where we plot the ratio between the MAP and truth pixel values in harmonic space.

We also construct and apply an HMC-NUT sampler similar to the previous section but with the additional amplitude dependence. The sampled spectra and the spectrum of the mean map are shown in Fig. 6 in shaded orange and in red respectively. Similar to the CMB model, the sampled spectra are unbiased. This is also confirmed by the parameter constraint as shown in Fig. 9, where we see a 2.9σ detection of the ISW amplitude. Meanwhile, as the amplitudes are now free, the spectrum of the mean map is still biased, but to lesser degrees than the MAP solution.

This is consistent with the field-level results (rows 4–5 of Fig. 8). Here, for the ISW tracer, the MAP map is essentially constant spatially, whereas the mean map still contains the right amount of power but has placed it all in the wrong place (as in the case of fixed amplitude).

V. JOINT RECONSTRUCTION OF THE CMB, ISW, AND THE GALAXY DENSITY MAPS

We now present the main analysis, where we generalize the framework to jointly analyze data from CMB and wide-field galaxy surveys on the field level. The main goal is the following. Given an observed temperature map and six tomographic galaxy density maps, we want to construct estimates of the primordial (s^P), the ISW (s^I), and the galaxy density ($\{s^{G,i}\}$) maps, along with two-point and cosmological parameter constraints, all in a consistent and computationally efficient Bayesian framework.

From now on, we will only consider the sampling approach. The addition of galaxy maps introduces off-diagonal terms in the covariance of the posterior distribution, as in Eq. (9). As we shall see, following the algorithmic prescription in Sec. II C, we can break the degeneracy between the primordial and the ISW maps using additional maps of the galaxy density.

A. Theoretical covariance

Equation (23) gives the general expression for the theoretical covariance of 2-dimensional tracer fields. In the case of CMB-ISW model, $\langle s^P s^I \rangle = 0$. Now, we

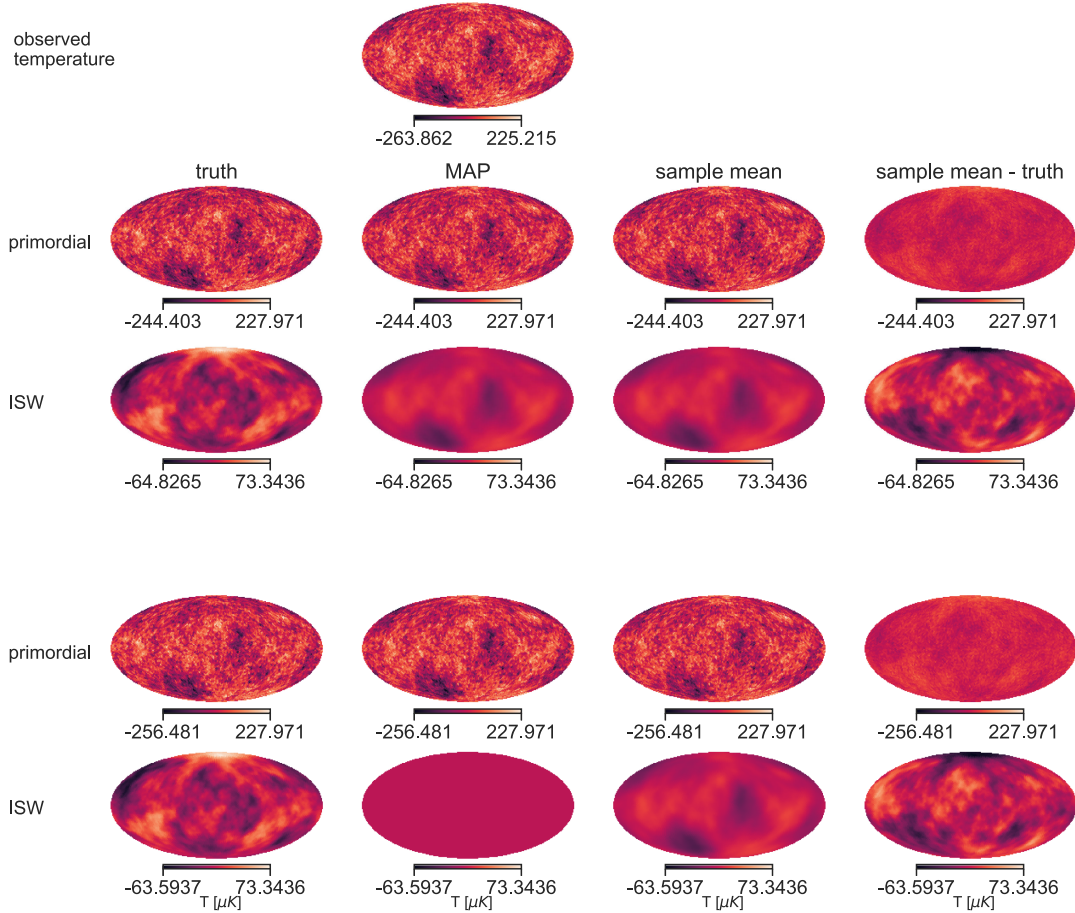


FIG. 8. Comparison between the observed data (row 1), truth maps, MAP maps, and the mean sampled maps. Rows 2–3 show the model with fixed amplitude, while rows 4–5 show the model of free amplitude. The color bar is shared across each row. Note that in the case of fixed amplitude, the mean sampled map is equal to the MAP map (second and third columns in the second and third rows), while when the amplitude is varied, the two diverge (same columns in the fourth and last row, but most obvious in the ISW maps in the last row) since the posterior marginalized over the amplitude is no longer a simple Gaussian in the signals.




FIG. 9. The posterior distribution of the primordial and ISW amplitudes with parameter bounds. The data strongly constrain the sum of the two amplitudes (a fact not obvious here because the x -range is so much smaller than the y -range), but the sampling technique manages to correctly infer the values of both.

introduce tomographic galaxy tracers, which correlate with the ISW effect (but not the primordial CMB) through their common dependence on the matter density field. This correlation will show up as off-diagonal terms in their covariance matrix, as indicated explicitly in Eq. (9).

Let $N_i(\chi)$ be the normalized line of sight galaxy density distribution for the redshift bin i , then the galaxy clustering window function that goes into Eq. (21) is given by

$$W^{g,i}(\chi) = b_i N_i(\chi), \quad (31)$$

where b_i is the linear galaxy bias that connects the matter power spectrum and the galaxy number density power spectrum. Throughout this study, we will treat each b_i as a scale-independent parameter with a fiducial value of 1.

For wide-field photometric surveys, the galaxy redshift distribution N_i varies widely between experiments and catalogs. We consider the MagLim sample from DES Year3 [46,47] as an example; projected distributions for LSST can be found in [48]. The MagLim catalog consists

TABLE I. The redshift parameters.

	bin 1	bin 2	bin 3	bin 4	bin 5	bin 6
z^c	0.30	0.47	0.62	0.78	0.90	1.00
z^w	0.10	0.07	0.07	0.07	0.05	0.05
Σ [deg ⁻²]	447.29	319.90	325.48	435.04	316.74	298.85

of 6 redshift bins spanning a redshift range of 0 to 1.2, calibrated using the self-organizing map methods (SOMPZ) and clustering redshifts. The catalog has been extensively tested on simulations and was used by the DES collaboration for the fiducial DES Year3 cosmology analysis [7,47,49]. For each redshift bin, we model $N_i(z)$ using the center (z_c) and width (z_w) of the distribution, following the functional form

$$\log N_i(z) \propto -\frac{1}{2} \left(\frac{z - z_i^c}{z_i^w} \right)^2. \quad (32)$$

We also used the number density (Σ_i) of the MagLim catalog for our simulations but assumed full-sky coverage instead of the DES footprint. We tabulate the binned z_c , z_w , and Σ in Table I and plot the normalized redshift distribution in Fig. 10.

Using the window function in Eq. (31), the covariance and the correlation coefficients between the ISW effect and the galaxy density can be computed using Eq. (23) and (24) (recall that primordial CMB is independent of the other tracers). The correlation coefficients are shown in the lower left corner of Fig. 11 in blue. This correlation is the information we hope to leverage to break the degeneracy between the primordial and the ISW maps. The bottom-left panel demonstrates that the ISW effect is most strongly correlated with the galaxy maps on large scales at low redshift.




FIG. 10. The galaxy number densities for each redshift bin as a function of redshift z . $N_i(z)$ is qualitatively modeled on the DES Year 3 MagLim sample.

B. Cosmological parameters

Let $\mathbb{C}^{\mathcal{P}}$, $\mathbb{C}^{\mathcal{I}}$, and $\mathbb{C}^{\mathcal{G},i}$ be the fiducial power spectra of the primordial, ISW, and the galaxy fields. Similar to the CMB-ISW model, we introduce an amplitude parameter for each tracer

$$\begin{aligned} \mathbb{C}^{\mathcal{P}} &\rightarrow (A^{\mathcal{P}})^2 \mathbb{C}^{\mathcal{P}}, \\ \mathbb{C}^{\mathcal{I}} &\rightarrow (A^{\mathcal{I}})^2 \mathbb{C}^{\mathcal{I}}. \end{aligned} \quad (33)$$

Equation (33) encodes a powerful stress test of Λ CDM. Consider the fiducial growth function $D(z)$. If the truth cosmology deviates from Λ CDM, then the actual growth function will be $A(z)D(z)$, where $A(z)$ is some redshift-dependent factor. Thus, we can interpret $A^{\mathcal{I}}$ as an integral of $A(z)$ over the ISW window function. Then, any detection of $A^{\mathcal{I}} \neq 1$ implies a deviation from the Λ CDM model. The ISW window function is rather wide.

The tomographic galaxy density window function is much narrower. Therefore, multiplying each binned galaxy spectrum by an amplitude factor $A^{\mathcal{G},i}$ would inform us of the consistency of Λ CDM at each redshift slice. This would be, and ultimately will be, a much more strenuous test of the Λ CDM model. Unfortunately, since we consider here only galaxy maps, $A^{\mathcal{G},i}$'s are entirely degenerate with the b_i 's in Eq. (31). So in this study, we will set $A^{\mathcal{G},i} = 1$. In future works, one could jointly analyze galaxy and shear maps to break this degeneracy and directly constrain the amplitude of the growth function in relatively narrow redshift intervals.

C. Noise model and simulation

Since we are mostly interested in extracting the large-scale ISW signal, we again perform the simulation on a HEALPix grid of $N_{\text{SIDE}} = 32$. We generate the temperature and galaxy tracers using the full covariance as described in Sec. VA. For the observed temperature map, we inject white noise with a variance of $200 \mu\text{K}^2$. For the tomographic galaxy maps, we assume a noise spectrum of

$$\mathbb{C}^{n,\mathcal{G},i}(l) = \frac{\pi^2}{180^2 \Sigma_i}, \quad (34)$$

where Σ_i is the number density of bin i per square degrees. The realized spectra (black) and the modeled noise spectra (purple) are shown in Fig. 12 (for the primordial and the ISW signals) and the upper right panels of Fig. 11 (for the ISW and the galaxy signals).

D. Posterior modeling and sampling

The observables are

$$d^{\mathcal{I}} = s^{\mathcal{P}} + s^{\mathcal{I}} + n^{\mathcal{I}}, \quad (35)$$




FIG. 11. The correlation structure between different tracers and the sampler results. The lower-left portion shows the correlation coefficients (blue) between the ISW and the 6 DES Year3 MagLim-like galaxy tracers from different redshift bins. The primordial field is not shown here since it is independent of all late-time tracers. One crucial observation is that (lowest-left panel) the ISW effect correlates most strongly with the low redshift galaxy density maps, since only at late times did dark energy significantly drive the accelerated cosmic expansion. The upper right corner shows the distribution of power spectra sampled from the joint posterior Eq. (37) compared to the true power spectra. The input galaxy and ISW signals are shown in black while the distribution of the posterior samples is shown in orange. The diagonal subplots also contain the noise (auto)spectra. Note that the galaxy autospectrum is unitless since the galaxy redshift distribution is normalized. The ISW autospectrum has the unit of μK^2 . Because of the different numeric scales, the galaxy-ISW cross-spectra and the ISW autospectrum are broken off into their own column with independent y-axis limits.

$$d^{\mathcal{G},i} = s^{\mathcal{G},i} + n^{\mathcal{G},i}. \quad (36)$$

Therefore, the inference problem is specified by the posterior distribution

$$p(s^{\mathcal{P}}, \{s_i^{\mathcal{G}}\}, s^{\mathcal{I}}, A^{\mathcal{P}}, A^{\mathcal{I}}, \{b_i\} | d^{\mathcal{I}}, \{d^{\mathcal{G},i}\}) \propto \left(\det \left[\mathbb{C}^{n,\mathcal{T}} \prod_{i=1}^6 \mathbb{C}^{n,\mathcal{G},i} \right] \det [\mathbb{C}_l(b_1, \dots, b_6, A^{\mathcal{I}}, A^{\mathcal{P}})] \right)^{-\frac{1}{2}} \\ \times \exp \sum_{lm} \frac{-|d_{lm}^{\mathcal{I}} - s_{lm}^{\mathcal{P}} - s_{lm}^{\mathcal{I}}|^2}{2\mathbb{C}_l^{n,\mathcal{T}}} \exp \sum_{i=1}^6 \sum_{lm} \frac{-|d_{lm}^{\mathcal{G},i} - s_{lm,i}^{\mathcal{G}}|^2}{2\mathbb{C}_l^{n,\mathcal{G},i}} \exp \sum_{lm} \frac{-|\{s_{lm,1}^{\mathcal{G}}, \dots, s_{lm,6}^{\mathcal{G}}, s_{lm}^{\mathcal{I}}, s_{lm}^{\mathcal{P}}\}|^2}{2\mathbb{C}_l(b_1, \dots, b_6, A^{\mathcal{I}}, A^{\mathcal{P}})}, \quad (37)$$

where \mathbb{C}_l is given in Eq. (9) and the notation $|x|^2/\mathbb{C}$ represents the quadratic form $x^T \mathbb{C}^{-1} x$




FIG. 12. The recovered $s^{\mathcal{P}}$ and $s^{\mathcal{I}}$ spectra using the joint reconstruction technique described in Sec. V. For each case, the truth spectrum is shown in black, the 1σ credible interval of the sampled spectra is shown in shaded orange, and the spectrum of the mean map is shown in red. The effective noise of the primordial maps is $\mathbb{C}^{\mathcal{I}} + \mathbb{C}^{n,\mathcal{I}}$ shown in purple. The effective noise of the CMB map ($\mathbb{C}^{\mathcal{P}} + \mathbb{C}^{n,\mathcal{I}}$) is above the ISW spectra on all scales and hence not shown in the plot. Notice that both tracers are reconstructed with much higher SNR comparing to the CMB-ISW model.

The first determinant terms encode the noise covariance, which is fixed in our example. The \mathbb{C}_i determinant encodes the posterior's dependence on the tracer amplitudes. The first two exponential terms come from the likelihood of the temperature and the galaxy maps. The last exponential term is the Gaussian priors on the signals, including the off-diagonal covariance matrix from Eq. (9). Not shown here is that all the free amplitudes have flat priors in the interval $[0.2, 3]$.

The structure of the sampler is much the same as in the CMB and the CMB-ISW models, following the prescription of Sec. II C. Notice that since now the covariance $\mathbb{C}_i(b_1, \dots, b_6, A^{\mathcal{I}}, A^{\mathcal{P}})$ is nondiagonal and extremely high dimensional, we have to employ the block diagonal Cholesky decomposition method introduced in Sec. II C to make the sampler computationally feasible.

Despite the high dimensionality of the problem, we find that the chain equilibrates quickly (more details are given in Appendix C). In general, the amplitude parameters have a much longer correlation length than latent map parameters during the sampling phase. Among the amplitude parameters, the ISW amplitude has a much longer correlation length of 10^2 samples. Overall, the sampler is very fast. The entire analysis took less than 10 hours on a single Apple M1 CPU.

E. Results

For each iteration of the sampler after burn-in, we collect a set of maps and parameters

$$\{\mathbf{s}_i, \mathbf{A}_i\} = \{s_{lm,i}^{\mathcal{P}}, s_{lm,i}^{\mathcal{I}}, \{s_{lm,k}^{\mathcal{G}}\}_i, A^{\mathcal{P}}, A^{\mathcal{I}}, \{b_k\}_i\} \quad (38)$$

for $k = 1, \dots, 6$ labeled by a common sample index i . The collection $\{\mathbf{s}_i\}$ forms the set of posterior samples which we will now analyze. We will present our findings in three parts; cosmological parameter constraints, power spectra reconstruction, and field-level reconstruction. The final results for this joint analysis are shown in Figs. 11–13.

1. Cosmological parameter constraints

The constraints on the two temperature tracer amplitudes and the six galaxy biases are shown in Fig. 13 and summarized in Table II. All 8 parameters are unbiased within 2σ . Consistent with our previous findings, the best-constrained parameters are the primordial amplitude and the tomographic galaxy biases which all have small effective noise. For these parameters, we achieve percent-level constraints assuming our very simple problem setups. We find that $A^{\mathcal{I}}$ is also unbiased and constrained to around 15% (or a 6.9σ detection), improving dramatically compared to the $\sim 40\%$ constraint (or 2.9σ detection) in the absence of galaxy data.

2. Power spectra constraints

Besides obtaining the correct overall power spectra amplitudes, we show that the reconstructed power spectra are unbiased for all tracers for all scales. The results for the primordial CMB and the ISW effect are shown in Fig. 12, where the sampling result is shown in shaded orange and the truth is shown in black. For ISW, we further observe that the uncertainty of the power spectrum estimation also shrinks considerably around the truth compared to the CMB-ISW model. This gain in SNR is directly attributed to the new information from the galaxy tracer fields.




FIG. 13. The posterior distribution of the amplitudes A^P and A^I and the six galaxy biases b_i 's. The 2-dimensional contours label the 0.68 and 0.95 credible intervals. The KDE-smoothed histograms above the contours are the parameters' marginal distributions, where the shaded intervals represent the 0.16 and 0.84 quantiles. The truth values are indicated by dashed black lines.

The sampling result for the ISW and the galaxy tracers (all the tracers that are correlated with each other) are shown in the upper right panels of Fig. 11 in orange. Here we see that the method has captured all the auto and cross-spectra of the tracer fields. The galaxy power spectra are especially well-reconstructed, in part due to their intrinsic high SNR observations.

3. Field-level reconstruction

The CMB-ISW-galaxy model reconstructs tracer maps at higher SNR than previous models. The result for the

temperature tracers is shown in Fig. 14. Comparing to the CMB-ISW model (Fig. 8), we see a dramatic improvement in the reconstruction accuracy. Under the multiprobe joint reconstruction framework, the field-level information in the galaxy maps funnels into the temperature map-making process and efficiently breaks the degeneracy between the primordial and the ISW field. Most notably, although the ISW signal is by far noise dominated on all scales, the mean sampled ISW field is now actually tracing the structures of the true ISW field and decorrelated with the true primordial field. The primordial reconstruction also

TABLE II. The cosmological parameter constraints given by the CMB-ISW-galaxy model. Here $A^{\mathcal{P}}$ and $A^{\mathcal{I}}$ are the amplitudes of the primordial and the ISW power spectra and b_i are the tomographic galaxy biases.

Parameter	Constraint
$A^{\mathcal{P}}$	$1.0029^{+0.0085}_{-0.0089}$
$A^{\mathcal{I}}$	$1.03^{+0.17}_{-0.15}$
b_1	$1.0000^{+0.0070}_{-0.0084}$
b_2	$0.9882^{+0.0083}_{-0.0080}$
b_3	$1.0068^{+0.0081}_{-0.0082}$
b_4	$0.9900^{+0.0083}_{-0.0074}$
b_5	$1.0052^{+0.0081}_{-0.0079}$
b_6	$1.0060^{+0.0080}_{-0.0094}$

receives the same benefit, as the residual error of its reconstruction is visibly less correlated with the true ISW field, compared to the CMB-ISW model.

The final result also includes samples of denoised tomographic galaxy maps as shown in Fig. 15. In theory, the cross-correlation with the ISW component of the temperature map can boost the SNR of the reconstructed galaxy maps as well. However, since the observed galaxy maps are already high in SNR and the degrees of freedom in the galaxy maps by far overwhelm that of the temperature tracers, their improvement is negligible.

VI. CONCLUSIONS

We have implemented a general hierarchical Bayesian framework that employs HMC to sample directly from the

joint posterior of the field-level multiprobe model. We did this in a pristine framework; simulated all-sky maps with simple noise properties. One goal of this is to understand the advantages and limitations of field-level analyses, prior to including more realistic effects. The other is more specific—to assess how accurately surveys can measure the ISW amplitude.

We particularly focused on comparing two approaches: maximum *a posteriori* (MAP) and sampling. This enabled us to demonstrate both the well-known bias of the MAP two-point estimator [37] (e.g., the lower jagged curves in Fig. 2) and the multiplicative bias of the field-level values (Fig. 3). These MAP biases persist as we added more complexity to the data vector and the contributing signals. The sampling approach is also biased at the field level (although, in general to a lesser extent) but is unbiased for power spectra and cosmological parameter constraints, as illustrated in Figs. 4, 9, and 13. This suggests that *the Bayesian posterior sampler produces unbiased cosmological parameters when multiple surveys are analyzed jointly*. Given the potential biases involved in map making and then cross-correlating, this seems to us to provide an excellent justification for the use of field-level analyses moving forward.

Our results addressing the second goal can be expressed in a single number. Using only CMB data, the amplitude of the ISW signal can be extracted from the CMB only with little power, perhaps at the $2\text{--}3\sigma$ level. However, when galaxy survey data is added, we project a 6.9σ detection. Before exploring the limitations of this projection, it is worth emphasizing that the long-term goal is to stress test Λ CDM, and this method provides a test at the 10–15%




FIG. 14. Comparison between the observed temperature (row 1), the primordial CMB maps (row 2), and the ISW maps (row 3). For both the primordial CMB and the ISW maps, we compare the truth, the mean sampled, and the difference maps obtained from the CMB-ISW-galaxy model with free cosmological parameters.




FIG. 15. Map-level reconstruction of the galaxy fields. Here we only show the first and the last redshift bins, the results for other bins are similar. Unit is in dimensionless overdensity. Note that although the two redshift bins have the same map resolution, the pixel scale of the lower redshift bin corresponds to smaller physical scale. Thus, the lower redshift bin has higher map-level variance, as expected.

level. However, this is but one of a slew of amplitudes that can be measured with upcoming data so we are optimistic about this general idea of constraining amplitudes of large-scale power spectra as a powerful way of testing the fiducial cosmological model.

Back to the limitations of our analysis; assuming full-sky slightly overstates the capability of the CMB, which is masked in the Galactic plane and overstates by at least a factor of two the coverage of, e.g., LSST. Statistically, then, one might reasonably inflate our projections by $\sqrt{2}$. However, there are a number of signals that we did *not* include: galaxy shapes and CMB lensing [25–27]. The kernels for both of these—especially the former—overlap significantly with that of ISW, so we expect that including them will quite likely recover this factor of $\sqrt{2}$. By adding these observables into the analysis, we could construct tomographic (convolved with different, albeit overlapping, kernels) maps of the matter density in a consistent Bayesian framework.

However, before turning to real data, we must relax the simplifications of our simulations, so we spend the rest of this conclusion alerting ourselves and our readers to those hurdles.

A. General cosmological parameters

There is no conceptual barrier to including cosmological parameters that modify the shape of the cross-spectra (in contrast to A and b which modulate only their amplitudes). The main challenge is that we must specify the derivatives of the posterior distributions with respect to each of the cosmological parameters in a computationally efficient fashion (e.g., finite difference methods will be too inefficient in an inference algorithm of this scale). The process of differentiating through the Boltzmann code and the Limber approximation computation is especially difficult. We see a few ways for future projects to tackle this issue:

- (i) Use automatic differentiation to take gradients through the cosmological dependence [50].
- (ii) Use simple fitting functions (e.g., [51]) where the analytical derivatives are easily attainable.
- (iii) Train a neural network-based cosmological emulator where the network is by definition differentiable [52,53].
- (iv) Exclude cosmological parameters from the HMC sampling altogether. Instead, we can sample the cosmological parameters and the maps iteratively in the Gibbs sampling paradigm [12,27,54].

B. Masking, anisotropic noise, and other systematic effects

The likelihood model we used in this study is very simple. We considered only the case of signal reconstruction on the full sky with isotropic noise and no masking. In order to adapt this algorithm for real-data analysis, we must take into account the limitation of survey geometry for both experiments, as well as foreground and point source masks. Further, the noise in real data will often be anisotropic and often a parametric function of a set of systematics spatial templates. Both the masking and the anisotropic noise models will introduce off-diagonal terms in the covariance matrix computation in Fourier space. Thus, a computationally efficient solution is to still sample the full sky, unmasked, and noiseless maps in Fourier space, transform the maps into real space, and define the likelihood there. Once in real space, we can implement different anisotropic and parametric noise models, and even attempt to constrain nuisance noise model parameters during sampling as well.

Systematic effects such as foregrounds and survey properties can be handled in the general framework of this field-level analysis. In particular, the *signal* vector can be expanded to include these. This one-step approach—as opposed to the current treatments—may be necessary for future surveys with increased statistical precision.

One simple way to understand why this may be needed is that a sample of cosmological parameters that predicts large clustering is more likely to label an ambiguous object a galaxy (rather than a star) if it is near another galaxy. Other systematics—such as photometric redshift uncertainty—can be included by introducing nuisance parameters.

C. Smaller scales

We have included only large scales here, and there is a huge advantage to doing so, in that the prior distributions of the signals is known to be Gaussian. There is a huge disadvantage to throwing out all the information available on small scales. Including small scales in the posterior requires a knowledge of the prior distribution of the signal, a distribution that is less and less Gaussian as we push to smaller scales. There are two possible approaches to this: (i) assume a simple distribution (e.g., Gaussian or log-normal [52]) and investigate the potential biases by running the pipeline on simulations; and (ii) the ambitious approach of rolling the clock back and using the primordial fields as the parameters given the observed highly processed fields (e.g., [55,56]).

ACKNOWLEDGMENTS

This work is supported by U.S. Dept. of Energy Contract No. DE-SC0019248 and by NSF Award No. 2020295. S. D. is grateful to the Aspen Center for Physics, where a workshop in Summer 2022 exposed him to some of these ideas. We are also grateful to Alan Heavens, Andrew Jaffe, Xiangchong Li, Marius Millea, Chirag Modi, Fabian Schmidt, and Ben Wandelt for useful conversations. We would also like to thank the anonymous reviewer for their constructive comments.

APPENDIX A: DISTRIBUTIONS OF INDIVIDUAL MODES OF THE SAMPLED POWER SPECTRA

Here we present the marginal posterior distribution of \mathcal{C} for each individual l mode for the CMB model ($\mathcal{C}^{\mathcal{P}}$, Fig. 16), CMB-ISW model ($\mathcal{C}^{\mathcal{I}}$, Fig. 17), and the CMB-ISW-galaxy model ($\mathcal{C}^{\mathcal{I}}$, Fig. 18).

Figures 16 and 17 demonstrate that the MAP power-spectra amplitudes (vertical blue and orange lines) have greater bias 1) at higher l -modes where the noise power is larger 2) when the power-spectra amplitudes are set free. Figures 17 and 18 show that, in the case of free amplitude parameters, the quality of $\mathcal{C}^{\mathcal{I}}$ reconstruction (the width of its marginal distribution) improves dramatically when one introduces galaxy information.

APPENDIX B: MAP AND SAMPLER ESTIMATORS

Our algorithm generates samples of the posterior. In the text we made two claims about these samples (for the case




FIG. 16. The marginal distribution of $\mathcal{C}^{\mathcal{P}}$ of the posterior samples for the CMB model (Sec. III). The case with fixed fiducial power spectrum is shown in orange, and the case where one admits a free amplitude is shown in blue. The vertical black lines show the truth $\mathcal{C}^{\mathcal{P}}$, while the colored vertical lines show the MAP values for $\mathcal{C}^{\mathcal{P}}$.




FIG. 17. The marginal distribution of $\mathcal{C}^{\mathcal{I}}$ of the posterior samples for the CMB-ISW model (Sec. IV). The case with fixed fiducial power spectrum is shown in orange, and the case where one admits a free amplitude is shown in blue. The vertical black lines show the truth $\mathcal{C}^{\mathcal{I}}$, while the colored vertical lines show the MAP values for $\mathcal{C}^{\mathcal{I}}$.




FIG. 18. The blue lines show the marginal distribution of $C_l^{\mathcal{I}}$ of the posterior samples for the CMB-ISW-galaxy model (Sec. V). The black and the red dashed lines mark the truth value and the mean of the sample distribution, respectively.

of fixed cosmological parameter): (i) the mean (two-point and n -point) power spectrum of the samples is unbiased; and (ii) in the single component case, the mean map is equal to the MAP map. To demonstrate this, we let $\{s_{lm,i}^\alpha\}$ be the set of maps that we have collected during sampling, where i denotes the sample index and α denotes the type of cosmological probe. Let N be the total number of samples.

1. Power spectrum

We first show that the mean sampled n -point power spectrum estimator is unbiased. Let us consider the pixel variance, which is the simplest two-point statistic. For a fixed data set D , the mean pixel variance across all the posterior samples is $\text{mean}(\langle s^2 \rangle)$. Here, s is the signal map vector, $\langle s^2 \rangle$ denotes the average pixel variance of s , and the “mean” operation is taken over all the posterior samples. Using i as the index for the posterior samples,

$$\text{mean}(\langle s^2 \rangle) \quad (\text{B1})$$

$$= \lim_{N \rightarrow \infty} \sum_{i=1}^N \frac{\langle s_i^2 \rangle}{N} \quad (\text{B2})$$

$$= \int ds |s^2| p(s|D), \quad (\text{B3})$$

where the short hand $|s^2|$ means $\sum_{\alpha=1, \dots, N_{\text{pix}}} s_\alpha^2 / N_{\text{pix}}$ with α being the pixel index. Equation (B3) in general depends on the data D . Practically, this means that $\text{mean}(\langle s^2 \rangle)$ is biased

only in the sense that the observed data has intrinsic randomness. If we proceed to integrate over the intrinsic variance in the data, we have

$$\int ds \int dD |s^2| p(s|D) p(D) \quad (\text{B4})$$

$$= \int ds |s^2| p(s) \quad (\text{B5})$$

$$= \langle s^2 \rangle \quad (\text{B6})$$

which is the true pixel variance the signal. This argument works for the covariance between two different pixels, and more generally, for the n -point correlation functions.

2. MAP estimator

We claim that the mean sampled map is the field-level MAP solution, and the power spectrum of the mean map is the MAP spectrum. Similar to the argument above, we start with

$$\langle s \rangle = \int ds s p(s|D), \quad (\text{B7})$$

where $p(s|D)$ is Gaussian. Since the mean of a Gaussian distribution is also the point of maximum probability, $\langle s \rangle$ is thus also the MAP estimator. More explicitly, by independence between different l, m -modes, we can suppress the l, m subscript on s and look at the term

$$\text{mean}(s_i) \times p(D) \quad (\text{B8})$$

$$= \lim_{N \rightarrow \infty} \sum_{i=1}^N \frac{s_i}{N} \times p(D) \quad (\text{B9})$$

$$= \int ds \frac{s}{2\pi\sqrt{C^n C^s}} \exp\left(-\frac{(D-s)^2}{2C^n} - \frac{s^2}{2C^s}\right) \quad (\text{B10})$$

$$= \frac{C^s D}{C^n + C^s} \times \frac{\exp(-\frac{D^2}{2(C^n + C^s)})}{(2\pi(C^n + C^s))^{1/2}}, \quad (\text{B11})$$

which is the Wiener solution for s times the probability of the data. Dividing by $p(D)$ yields

$$\langle s \rangle = \frac{C^s D}{C^n + C^s} \quad (\text{B12})$$

as desired. It follows that the power spectrum of the mean map is the Wiener filter spectrum.

Interestingly, notice that the above arguments rests on the symmetry of the Gaussian distribution. If we introduce an additional amplitude parameter A , the distribution over A is

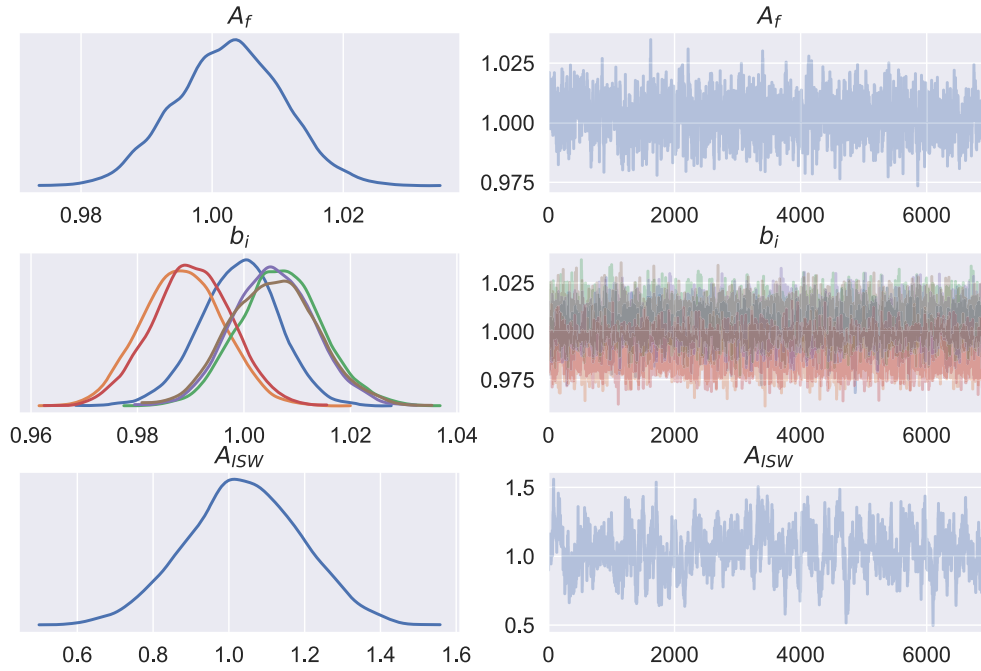


FIG. 19. The chain convergence of the CMB-ISW-galaxy model for the 8 sampled cosmological parameters (the power spectra amplitudes and galaxy biases). The left panels show the posterior distribution of each parameter after KDE smoothing. The right panels show each parameter’s value as a function of sample index after 3000 warmup steps.

no longer Gaussian, and so the mean map is no longer the MAP map solution (as discussed at length in the main text).

APPENDIX C: CHAIN CONVERGENCE

Here we present the chain convergence information of the CMB-ISW-galaxy model, focusing on the eight cosmological parameters. The sampled parameter values as a

function of sample index is shown in Fig. 19, and the corresponding autocorrelation functions are shown in Fig. 20. In general, we observe that the map parameters (not shown here) have much shorter correlation lengths than the cosmological parameters. Among the cosmological parameters, A^P and b_i have much shorter correlation lengths compared to A^I .




FIG. 20. The autocorrelation functions of the eight sampled cosmological parameters shown in Fig. 19.

- [1] N. Aghanim *et al.* (Planck Collaboration), *Astron. Astrophys.* **641**, A1 (2020).
- [2] N. Aghanim *et al.* (Planck Collaboration), *Astron. Astrophys.* **641**, A8 (2020).
- [3] S. Aiola *et al.*, *J. Cosmol. Astropart. Phys.* **12** (2020) 047.
- [4] O. Darwish *et al.*, *Mon. Not. R. Astron. Soc.* **500**, 2250 (2020).
- [5] F. Bianchini *et al.*, *Astrophys. J.* **888**, 119 (2020).
- [6] L. Balkenhol *et al.*, *Phys. Rev. D* **108**, 023510 (2023).
- [7] T. M. C. Abbott *et al.* (DES Collaboration), *Phys. Rev. D* **105**, 023520 (2022).
- [8] C. Heymans *et al.*, *Astron. Astrophys.* **646**, A140 (2021).
- [9] T. Hamana *et al.*, *Publ. Astron. Soc. Jpn.* **72**, 488 (2020).
- [10] T. M. C. Abbott *et al.*, *Phys. Rev. D* **107**, 023531 (2023).
- [11] M. Tegmark, *Astrophys. J.* **480**, L87 (1997).
- [12] B. D. Wandelt, D. L. Larson, and A. Lakshminarayanan, *Phys. Rev. D* **70**, 083511 (2004).
- [13] D. L. Larson, H. K. Eriksen, B. D. Wandelt, K. M. Gorski, G. Huey, J. B. Jewell, and I. J. O'Dwyer, *Astrophys. J.* **656**, 653 (2007).
- [14] H. K. Eriksen, J. B. Jewell, C. Dickinson, A. J. Banday, K. M. Gorski, and C. R. Lawrence, *Astrophys. J.* **676**, 10 (2008).
- [15] D. L. Larson, H. K. Eriksen, B. D. Wandelt, K. M. Gorski, G. Huey, J. B. Jewell, and I. J. O'Dwyer, *Astrophys. J.* **656**, 653 (2007).
- [16] E. Anderes, B. D. Wandelt, and G. Lavaux, *Astrophys. J.* **808**, 152 (2015).
- [17] N. Aghanim *et al.* (Planck Collaboration), *Astron. Astrophys.* **641**, A5 (2020).
- [18] E. Tsaprazi, N.-M. Nguyen, J. Jasche, F. Schmidt, and G. Lavaux, *J. Cosmol. Astropart. Phys.* **08** (2022) 003.
- [19] A. Loureiro, L. Whiteway, E. Sellentin, J. S. Lafaerie, A. H. Jaffe, and A. F. Heavens, *arXiv:2210.13260*.
- [20] A. Andrews, J. Jasche, G. Lavaux, and F. Schmidt, *Mon. Not. R. Astron. Soc.* **520**, 5746 (2023).
- [21] P. Fiedorowicz, E. Rozo, and S. S. Boruah, *arXiv:2210.12280*.
- [22] J. Jasche and G. Lavaux, *Mon. Not. R. Astron. Soc.* **447**, 1204 (2015).
- [23] D. K. Ramanah, G. Lavaux, J. Jasche, and B. D. Wandelt, *Astron. Astrophys.* **621**, A69 (2019).
- [24] W. Hu, *Astrophys. J. Lett.* **557**, L79 (2001).
- [25] M. Millea, E. Anderes, and B. D. Wandelt, *Phys. Rev. D* **100**, 023509 (2019).
- [26] M. Millea, E. Anderes, and B. D. Wandelt, *Phys. Rev. D* **102**, 123542 (2020).
- [27] M. Millea *et al.*, *Astrophys. J.* **922**, 259 (2021).
- [28] Q. Hang, S. Alam, J. A. Peacock, and Y.-C. Cai, *Mon. Not. R. Astron. Soc.* **501**, 1481 (2021).
- [29] K. M. Gorski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelman, *Astrophys. J.* **622**, 759 (2005).
- [30] A. Zonca, L. Singer, D. Lenz, M. Reinecke, C. Rosset, E. Hivon, and K. Gorski, *J. Open Source Software* **4**, 1298 (2019).
- [31] M. Betancourt, *arXiv:1701.02434*.
- [32] R. M. Neal, in *Handbook of Markov Chain Monte Carlo*, edited by S. Brooks, A. Gelman, G. Jones, and X.-L. Meng (Chapman & Hall/CRC, 2011), pp. 113–162, <https://doi.org/10.1201%2Fb10905>.
- [33] M. D. Hoffman and A. Gelman, *arXiv:1111.4246*.
- [34] E. Bingham, J. P. Chen, M. Jankowiak, F. Obermeyer, N. Pradhan, T. Karaletsos, R. Singh, P. Szerlip, P. Horsfall, and N. D. Goodman, *arXiv:1810.09538*.
- [35] D. Phan, N. Pradhan, and M. Jankowiak, *arXiv:1912.11554*.
- [36] Y. Akrami *et al.* (Planck Collaboration), *Astron. Astrophys.* **641**, A9 (2020).
- [37] G. B. Rybicki and W. H. Press, *Astrophys. J.* **398**, 169 (1992).
- [38] A. Kostic, N.-M. Nguyen, F. Schmidt, and M. Reinecke, *J. Cosmol. Astropart. Phys.* **07** (2023) 063.
- [39] L. A. Kofman and A. A. Starobinskii, *Sov. Astron. Lett.* **11**, 271 (1985).
- [40] R. G. Crittenden and N. Turok, *Phys. Rev. Lett.* **76**, 575 (1996).
- [41] P. Fosalba and E. Gaztanaga, *Mon. Not. R. Astron. Soc.* **350**, L37 (2004).
- [42] S. Boughn and R. Crittenden, *Nature (London)* **427**, 45 (2004).
- [43] R. Scranton, A. J. Connolly, R. C. Nichol, A. Stebbins, I. Szapudi, D. J. Eisenstein, N. Afshordi, T. Budavari, I. Csabai, J. A. Frieman, J. E. Gunn, D. Johnston, Y. Loh, R. H. Lupton, C. J. Miller, E. S. Sheldon, R. S. Sheth, A. S. Szalay, M. Tegmark, and Y. Xu, *arXiv:astro-ph/0307335*.
- [44] N. Afshordi, *Phys. Rev. D* **70**, 083536 (2004).
- [45] A. Manzotti and S. Dodelson, *Phys. Rev. D* **90**, 123009 (2014).
- [46] R. Cawthon *et al.* (DES Collaboration), *Mon. Not. R. Astron. Soc.* **513**, 5517 (2022).
- [47] G. Giannini *et al.*, *arXiv:2209.05853*.
- [48] S. S. Boruah, T. Eifler, V. Miranda, and P. M. S. Krishanth, *Mon. Not. R. Astron. Soc.* **518**, 4818 (2022).
- [49] M. Rodríguez-Monroy *et al.*, *Mon. Not. R. Astron. Soc.* **511**, 2665 (2022).
- [50] J.-E. Campagne, F. Lanusse, J. Zuntz, A. Boucaud, S. Casas, M. Karamanis, D. Kirkby, D. Lanzieri, Y. Li, and A. Peel, *Open J. Astrophys.* **6**, 1 (2023).
- [51] D. J. Eisenstein and W. Hu, *Astrophys. J.* **511**, 5 (1999).
- [52] S. S. Boruah, E. Rozo, and P. Fiedorowicz, *Mon. Not. R. Astron. Soc.* **516**, 4111 (2022).
- [53] T. Nishimichi, M. Takada, R. Takahashi, K. Osato, M. Shirasaki, T. Oogi, H. Miyatake, M. Oguri, R. Murata, Y. Kobayashi, and N. Yoshida, *Astrophys. J.* **884**, 29 (2019).
- [54] J. Alsing, A. Heavens, and A. H. Jaffe, *Mon. Not. R. Astron. Soc.* **466**, 3272 (2016).
- [55] A. E. Bayer, C. Modi, and S. Ferraro, *J. Cosmol. Astropart. Phys.* **06** (2023) 046.
- [56] C. Modi, F. Lanusse, U. Seljak, D. N. Spergel, and L. Perreault-Levasseur, *arXiv:2104.12864*.