

Constraining cosmological parameters with needlet internal linear combination maps. II. Likelihood-free inference on needlet internal linear combination power spectra

Kristen M. Surrao^{*} and J. Colin Hill[†]

Department of Physics, Columbia University, New York, New York 10027, USA

 (Received 25 June 2024; accepted 13 August 2024; published 4 September 2024)

Standard cosmic microwave background (CMB) analyses constrain cosmological and astrophysical parameters by fitting parametric models to multifrequency power spectra (MFPS). However, such methods do not optimally weight maps in power spectrum (PS) measurements for non-Gaussian cosmic microwave background (CMB) foregrounds. We propose needlet internal linear combination (NILC), operating on wavelets with compact support in pixel and harmonic space, as a weighting scheme to yield more optimal parameter constraints. In a companion paper, we derived an analytic formula for the NILC map PS, which is physically insightful but computationally difficult to use in parameter inference pipelines. In this work, we analytically show that fitting parametric templates to MFPS and the harmonic ILC PS yields identical parameter constraints when the number of sky components equals or exceeds the number of frequency channels. We numerically show that, for Gaussian random fields, the same holds for the NILC PS. This suggests that NILC can reduce parameter error bars in the presence of non-Gaussian fields since it uses non-Gaussian information. As Gaussian likelihoods may be inaccurate, we use likelihood-free inference with neural posterior estimation. We show that performing inference with autospectra and cross-PS of NILC component maps as summary statistics yields smaller parameter error bars than inference with MFPS. For a model with CMB, an amplified thermal Sunyaev-Zel'dovich (tSZ) signal, and noise, we find a 60% reduction in the area of the 2D 68% confidence region for component amplitude parameters inferred from the NILC PS, as compared to inference from MFPS. Primordial B -mode searches are a promising application for our new method, as the amplitude of the non-Gaussian dust foreground is known to be larger than a potential signal.

DOI: [10.1103/PhysRevD.110.063510](https://doi.org/10.1103/PhysRevD.110.063510)

I. INTRODUCTION

Standard state-of-the-art cosmic microwave background (CMB) data analyses—such as those used in *Planck* [1], the Atacama Cosmology Telescope (ACT) [2,3], the South Pole Telescope (SPT) [4], and BICEP [5]—fit data to theoretical parametric templates at the power spectrum level using multifrequency data. These analyses have been used to derive the parameter constraints that underlie the current standard model of cosmology. Before computing the frequency-frequency autospectra and cross spectra of the maps, inverse noise variance weighting is applied to the maps to downweight areas of high instrumental noise, allowing one to more optimally recover the signal (e.g., [6–9]). However, this weighting does not downweight CMB foreground contaminants, which are generally non-Gaussian fields. Therefore, the signal that is recovered when applying inverse noise variance weighting is the total

sky signal (foregrounds in addition to the CMB). Since the cosmological signal of interest is the CMB alone, we should aim to infer this particular component with the smallest possible error bars. In order to achieve that goal, one should apply a weighting scheme that captures both scale-dependent and spatially varying information to suppress the non-Gaussian CMB foregrounds, improving power spectrum estimation and hence parameter error bars. Since the weighting would be applied at the map level prior to power spectrum estimation, this motivates the exploration of various map-based weighting schemes.

At the map level, there are several methods for constructing component-separated maps of various signals in the microwave sky. Some methods fit parametric models [10,11]. Others, such as internal linear combination (ILC) [12–14], are semiblind approaches. The ILC procedure estimates a map of some signal of interest by finding the minimum-variance linear combination of the observed frequency maps that satisfies the constraint of unit response to the signal of interest. ILC can be performed in both real/pixel space and harmonic space.

^{*}Contact author: k.surrao@columbia.edu

[†]Contact author: jch2200@columbia.edu

Additionally, it can be applied on a tight frame of spherical wavelets, known as needlets, in a method called needlet ILC (NILC) [15]. The benefit of NILC is that needlets have compact support in both pixel and harmonic space, allowing the ILC weights to vary as functions of both angular scale and position. NILC has been used to construct high-resolution maps of individual components in the microwave sky using data or simulations from *Planck*, ACT, SPT, the Simons Observatory, CMB-S4, and *LiteBIRD*, e.g., Refs. [16–23].

The primary idea of this paper is to use NILC as a concrete example of a weighting scheme that optimally downweights contaminants and upweights signal-dominated regions before computing power spectra. In particular, we explore the approach of producing a NILC map for every component in the sky model, computing the autospectra and cross-power spectra of the resulting NILC maps, and using these power spectra as summary statistics for parameter inference. We do the same with harmonic ILC (HILC), mathematically proving that it yields the same parameter constraints as the traditional multifrequency power spectrum approach. This suggests that NILC can indeed reduce parameter error bars in the presence of non-Gaussian foregrounds since it uses non-Gaussian information that HILC does not. In a companion paper we developed an analytic formalism by which to parametrize NILC power spectra, which enabled novel insights into the beyond-Gaussian information that the NILC algorithm uses in constructing foreground-cleaned maps [24]. However, we found that parameter dependence of the NILC component map power spectra was complicated, particularly due to the appearance of terms involving correlations between the NILC weight maps and the underlying fields, and that the analytic results were very slow to use in a realistic computation.

In this paper, we investigate numerical techniques for parametrizing NILC map power spectra for use in a Gaussian likelihood. However, we also find that the Gaussian likelihood is not a good assumption, and thus use likelihood-free inference (LFI), showing explicitly how discrepant the results from the Gaussian likelihood are. Previous studies have considered the use of likelihood-free inference, also known as simulation-based inference (SBI), to constrain cosmological parameters (e.g., [25–27]). LFI can be performed at the field level, which indeed takes into account spatially varying information, but requires large numbers of computationally expensive simulations. It can also be done using summary statistics such as the power spectrum, which allows one to optimally compress the data for Gaussian random fields. Using NILC power spectra as summary statistics allows several advantages as compared to a pure field-level approach. First, a power-spectrum-based approach allows one to use noise-debiased data-split cross spectra (e.g., [28–30]), mitigating the need for the LFI to learn how to remove instrumental noise. Additionally,

nearly all the information for the primary field of interest (the CMB) is contained in its two-point function alone. NILC allows one to better recover that power spectrum by suppressing contributions from non-Gaussian contaminants.

The remainder of this paper is organized as follows. In Sec. II we review ILC methods, including harmonic and needlet ILC. In Sec. III we review likelihood-free inference (LFI), also known as simulation-based inference (SBI). In Sec. IV we provide an overview of the computational setup and the various methods implemented in this work. In Secs. V–VII, we describe in detail the various methods and summary statistics considered, using both Gaussian likelihoods and LFI to obtain posteriors. Specifically, we consider the use of multifrequency autospectra and cross-power spectra, harmonic ILC map power spectra, and needlet ILC map power spectra as summary statistics, respectively. In Sec. IX we present the results of each of the methods, showing how 2D posteriors shrink significantly when using needlet ILC summary statistics. Finally, in Sec. X we discuss the results and their implications, particularly for primordial *B*-mode searches. The main results of this work are summarized in Fig. 5 and Table II, as well as in Fig. 13 and Table V in Appendix D. Our code is publicly available in NILC-Inference-Pipeline [31] and contains pipelines for all the methods described in this work.

II. INTERNAL LINEAR COMBINATION

Internal linear combination (ILC) [12–15] is a method to estimate a map of a signal of interest by finding the minimum-variance linear combination of a set of observed frequency maps that satisfies the constraint of unit response to the signal of interest. Specifically, in pixel space, the signal of interest at some pixel p can be expressed as $\hat{y}(p) = \sum_i w_i \Delta T_i(p)$, where $\Delta T_i(p)$ is the temperature fluctuation in the i th frequency map and w_i is the associated linear-combination weight. Then the problem of finding the weights can be expressed as follows:

$$\min_{w^i} \sigma_{\hat{y}}^2 = N_{\text{pix}}^{-1} \sum_p (\hat{y}(p) - \langle y \rangle)^2$$

$$\text{such that } \sum_i w^i g_i = 1, \quad (1)$$

where N_{pix} is the number of pixels, $\langle y \rangle$ is the average signal across pixels, and g_i is the spectral response of the component of interest at the i th frequency channel. The solution for the weights can be found using Lagrange multipliers [14]:

$$w^i = \frac{g_j (\hat{R}^{-1})_{ij}}{g_k (\hat{R}^{-1})_{kl} g_l} \quad \text{with} \quad \hat{R}_{ij} = N_{\text{pix}}^{-1} \sum_p \Delta T_i(p) \Delta T_j(p), \quad (2)$$

where \hat{R}_{ij} is the empirical frequency-frequency covariance matrix of the observed maps. ILC can also be formulated in harmonic space, giving ℓ -dependent weights:

$$w_\ell^i = \frac{(\hat{R}_\ell^{-1})_{ij} g_j}{(\hat{R}_\ell^{-1})_{km} g_k g_m} \quad \text{with} \quad (\hat{R}_\ell)_{ij} = \sum_{\ell'=\ell-\Delta\ell/2}^{\ell+\Delta\ell/2} \frac{2\ell'+1}{4\pi} C_{\ell'}^{ij}. \quad (3)$$

The multipole bin width $\Delta\ell$ (or the pixel domain size in real-space ILC above) must be large enough to mitigate the ‘‘ILC bias’’ that results from computing the covariances for ILC weights using a small number of modes [15].

To maximize the robustness of the ILC procedure, it is frequently applied to CMB data on a needlet frame [1,15–18,32–36]. Needlets are a set of basis functions on the sphere that possess compact support in both real space and harmonic space [37,38], allowing one to obtain ILC weights that vary both as a function of scale (depending on multipole ℓ) and of position (depending on direction or spatial pixel \hat{n}). Such a procedure allows us to apply an optimal weighting scheme to non-Gaussian and/or anisotropic foregrounds. We briefly summarize the NILC procedure below [15,39].

Consider a set of frequency maps $T^i(\hat{n})$, where i is an index denoting the frequency channel of each map in our dataset. For a needlet filter $h_\ell^{(n)}$, indexed by (n) ranging from 1 to N_{scales} (the total number of needlet filters), the NILC operations on this set of frequency maps are as follows:

- (i) Transform each frequency map to harmonic space and filter it with the needlet filter function $h_\ell^{(n)}$:

$$T_{\ell m}^i \rightarrow T_{\ell m}^{i,(n)} \equiv T_{\ell m}^i h_\ell^{(n)}. \quad (4)$$

This procedure produces N_{scales} separate maps for each frequency channel (each frequency map gets filtered by each needlet filter separately).

- (ii) Define local pixel-space domains, and compute the smoothed frequency-frequency covariance matrix on each domain. In particular, let $\mathcal{D}_\alpha^{(n)}$ denote a real-space domain in frequency maps that have been filtered with needlet scale (n) , where α labels each domain on that map. The frequency-frequency covariance matrix is then

$$(\hat{R}_\alpha^{(n)})_{ij} = N_{\text{pix}}^{-1} \sum_{p \in \mathcal{D}_\alpha^{(n)}} \Delta T_i(p) \Delta T_j(p), \quad (5)$$

where N_{pix} is the number of pixels in $\mathcal{D}_\alpha^{(n)}$. This equation is nearly identical to Eq. (2) for the real-space ILC frequency-frequency covariance matrix, except that here the covariance matrix is

computed independently for each needlet filter scale. In practice, Eq. (5) is usually implemented by smoothing the product map $\Delta T_i \Delta T_j$ with a Gaussian kernel, with larger kernels used for lower-multipole needlet filter scales.

- (iii) For each frequency channel and each needlet filter scale, determine a map of weights in pixel space $W^{i,(n)}(\hat{n})$ analogously to Eq. (2). The weight maps are determined via the ILC algorithm, similar to Eq. (2) but performed on the local pixel-space domains from the previous step. Multiply each filtered frequency map by the associated weight map:

$$T^{i,(n)}(\hat{n}) \rightarrow \tilde{T}^{i,(n)}(\hat{n}) \equiv T^{i,(n)}(\hat{n}) W^{i,(n)}(\hat{n}). \quad (6)$$

- (iv) Add up the ILC-weighted maps to obtain a single ILC map at each needlet filter scale:

$$T^{\text{NILC},(n)}(\hat{n}) = \sum_i \tilde{T}^{i,(n)}(\hat{n}). \quad (7)$$

- (v) Apply the needlet filter $h_\ell^{(n)}$ again:

$$T_{\ell m}^{\text{NILC},(n)} \rightarrow T_{\ell m}^{\text{NILC},(n),(n)} \equiv T_{\ell m}^{\text{NILC},(n)} h_\ell^{(n)}. \quad (8)$$

- (vi) Sum the results from all needlet filter scales to obtain the final NILC map:

$$T^{\text{NILC}}(\hat{n}) = \sum_{(n)} T^{\text{NILC},(n),(n)}(\hat{n}). \quad (9)$$

Producing a NILC map in this manner allows the signal of interest to propagate in an unbiased fashion to the final map, due to the ILC signal-preservation constraint and the NILC filter power-preservation constraint, $\sum_{(n)} (h_\ell^{(n)})^2 = 1$ at each ℓ . However, contaminant signals propagate in a nontrivial way, as derived in Ref. [24].

Thus far, we have described ILC methods that remove contaminants by minimizing the variance of the final map. It is also possible to perform constrained ILC methods, which explicitly deproject contaminants by imposing an additional constraint that the response of the final map to the contaminant must be zero [40,41]. However, the additional constraint comes at the cost of increasing the variance of the final map. In this work, we consider only standard ILC methods (no explicit deprojection of contaminants) throughout.

III. LIKELIHOOD-FREE INFERENCE

In standard Bayesian analysis, given some parameters θ and an observation x , one can compute the posterior $p(\theta|x)$ using

$$p(\theta|x) = \frac{p(x|\theta)p(\theta)}{\int d\theta' p(x|\theta')p(\theta')}, \quad (10)$$

where $p(x|\theta)$ is the likelihood, $p(\theta)$ is the prior, and $p(x) = \int d\theta' p(x|\theta')p(\theta')$ is the evidence.

In certain situations, however, it is intractable to construct an analytic likelihood. Likelihood-free inference (LFI), also known as simulation-based inference (SBI) or implicit likelihood inference, is a method for determining posterior distributions without using an explicit likelihood (e.g., [42]). In particular, given some prior $p(\theta)$, LFI requires a simulator that samples from the prior and generates simulations, i.e., $x_i \sim p(x_i|\theta_i)$. With several simulations on hand, one can then train a normalizing flow network to learn the likelihood (“neural likelihood estimation”) [43–45], to learn the likelihood-evidence ratio (“neural ratio estimation”) [46–49], or to learn the posterior directly (“neural posterior estimation”) [50–53]. When the length of the data vector being simulated is much larger than the number of parameters, as is the case in our setup discussed in the coming sections, neural posterior estimation (NPE) generally performs best.

A. Masked autoregressive flows

There are various possible choices for the exact network that is used in LFI. In our work, we use masked autoregressive flows (MAF) [54], for which we find the most stable training results. MAF combines the advantages of autoregressive models [55], which model several conditionals whose product comprises the target joint density, and normalizing flows [56], which use invertible transformations to convert a base density into the target density. In the remainder of this section, we follow the notation of Ref. [54].

Autoregressive models decompose the target density $p(\mathbf{x})$ as

$$p(\mathbf{x}) = \prod_i p(x_i|\mathbf{x}_{1:i-1}), \quad (11)$$

where the subscript $1:i-1$ denotes that the i th state is dependent on all previous states indexed from 1 to $i-1$. The parameters of each conditional can be learned as functions of hidden states of recurrent networks.

Normalizing flows take some base density $\pi_u(\mathbf{u})$, draw $\mathbf{u} \sim \pi_u(\mathbf{u})$, and let $\mathbf{x} = f(\mathbf{u})$, where f is an invertible transformation. The joint target density $p(\mathbf{x})$ is computed as

$$p(\mathbf{x}) = \pi_u(f^{-1}(\mathbf{x})) \left| \det \left(\frac{\partial f^{-1}}{\partial \mathbf{x}} \right) \right|, \quad (12)$$

and thus, f must have a tractable Jacobian.

Autoregressive models can be treated as normalizing flows: Let

$$p(x_i|\mathbf{x}_{1:i-1}) = \mathcal{N}(x_i|\mu_i, (\exp \alpha_i)^2) \quad (13)$$

with $\mu_i = f_{\mu_i}(\mathbf{x}_{1:i-1})$ and $\alpha_i = f_{\alpha_i}(\mathbf{x}_{1:i-1})$, where f_{μ_i} and f_{α_i} are scalar functions. Then taking some random vector \mathbf{u} with $u_i \sim \mathcal{N}(0, 1)$, $x_i = u_i \exp \alpha_i + \mu_i$. This allows one to start with some random seed and transform it into the target density.

MAF is made by stacking several of the masked autoencoder(s) for distribution estimation (MADE) [57]. MADE uses a feedforward network to learn f_{μ_i} and f_{α_i} , applying binary masks to the weight matrices to enforce the autoregressive property. We refer the reader to Ref. [54] for more details.

IV. OVERVIEW OF METHODS AND COMPUTATIONAL SETUP

A. General computational setup

As a simple, demonstrative example, we consider a simulated sky model comprising two components: the CMB signal and the thermal Sunyaev-Zel’dovich (tSZ) signal, which is amplified by a factor of 150 at the map level (for reasons described below). The tSZ effect is the inverse-Compton scattering of CMB photons off hot electrons and is a useful probe of hot gas in galaxy clusters [58,59]. We refer to the amplified tSZ signal as the “fake” tSZ or “ftSZ” field. In CMB thermodynamic temperature units, the CMB spectral response is unity at every frequency. We use units of K for the CMB field. The tSZ spectral response, which we denote as $g(\nu)$ with ν being some frequency in Hz, is given by

$$g(\nu) = x \coth \frac{x}{2} - 4 \quad \text{with} \quad x = \frac{h\nu}{k_B T_{\text{CMB}}}, \quad (14)$$

where h is Planck’s constant, k_B is the Boltzmann constant, and $T_{\text{CMB}} = 2.726$ is the CMB temperature today (at redshift $z = 0$).

We simulate these sky components at two frequencies, 90 and 150 GHz. We then add Gaussian, white noise to the mock sky map for each frequency channel. Lensed $a_{\ell m}$ for the CMB component are obtained from the WebSky Extragalactic CMB Mocks,¹ which assume a flat Λ CDM cosmology with cosmological parameters consistent with Planck 2018 [29]: $(\Omega_m, \Omega_b, \sigma_8, n_s, h, \tau) = (0.31, 0.049, 0.81, 0.965, 0.68, 0.055)$ [60]. The HEALPIX/HEALPY [61,62] software package is used to compute some fiducial power spectrum and then generate many CMB realizations from that spectrum. For the tSZ maps, we use HALOSKY,² which generates realistic random non-Gaussian tSZ simulations from random halo catalogs. Specifically, HALOSKY

¹https://mocks.cita.utoronto.ca/index.php/WebSky_Extragalactic_CMB_Mocks.

²<https://github.com/marcelo-alvarez/halosky>.

Poisson samples from the Tinker *et al.* halo mass function [63] (determining halo abundance by using the WebSky [60] linear matter power spectrum based on the Planck 2018 cosmological parameters [29]) and then populates a catalog of halos along the light cone. The Battaglia *et al.* AGN feedback pressure profile is used [64]. We use redshift limits $0 \leq z \leq 5.0$ and mass limits $5 \times 10^{14} M_\odot \leq M \leq 10^{16} M_\odot$ in this construction, the latter chosen for computational efficiency. This narrow mass range is not important since we are rescaling the tSZ field by a large amplitude in this toy model. It also has the effect of populating fewer halos, making the tSZ effect more Poissonian and thereby more non-Gaussian. This is useful for our demonstration purposes, the reasoning for which is discussed further below.

To avoid noise biases in the measured power spectra, we use data-split cross spectra, as is done in many actual CMB analyses (e.g., [28–30]). For the noise power spectrum in each simulated split map, we use the model given by Ref. [65]:

$$N_\ell = W^2 e^{\ell(\ell+1)\sigma^2} \quad \text{with} \quad \sigma = \theta_{\text{FWHM}} / \sqrt{8 \ln 2},$$

where we set $W_{90} = W_{150} = 3 \times 10^4 \mu\text{K}_{\text{CMB}} \cdot \text{arcmin}$ for each noise split and $\theta_{\text{FWHM}} = 1.4 \text{ arcmin}$ for both the 90 and 150 GHz beams in our simulations. The large noise is important for our simple two-frequency, two-component-plus-noise sky model, where we do not want the NILC maps to completely clean the contaminating foregrounds (thus giving a realistic representation on what would happen with actual data where there are several sky components that cannot all be simultaneously cleaned completely). All maps are generated at HEALPIX resolution parameter $N_{\text{side}} = 128$, and power spectra are computed up to $\ell_{\text{max}} = 250$ for computational efficiency on this simple, demonstrative example. We then add the CMB, ftSZ, and noise map at each frequency and each split, resulting in four maps (two frequencies and two splits). An example of the simulated component power spectra and frequency-frequency power spectra is shown in Fig. 1.

Our goal is to constrain two parameters, A_{CMB} and A_{ftSZ} , which are two overall amplitude parameters that uniformly scale the CMB and ftSZ power spectrum over the entire multipole range, respectively. The fiducial parameter values are $A_{\text{CMB}} = 1 = A_{\text{ftSZ}}$. Specifically,

$$C_\ell^{ij}(A_{\text{CMB}}, A_{\text{ftSZ}}) = A_{\text{CMB}} C_\ell^{\text{CMB}} + A_{\text{ftSZ}} g_y^i g_y^j C_\ell^{\text{ftSZ}}, \quad (15)$$

where i and j index frequency channels, C_ℓ^{ij} is the frequency-frequency autospectrum or cross-power spectrum, C_ℓ^{CMB} is a CMB power spectrum template, C_ℓ^{ftSZ} is a power spectrum template of the amplified Compton- y field, and g_y^i and g_y^j denote the tSZ spectral responses at frequencies i and j . We could have also defined the




FIG. 1. Sample simulated power spectra (in units of K^2), plotted as $D_\ell = \ell(\ell + 1)C_\ell/(2\pi)$. Plotted are the CMB power spectrum (black dash-dotted curve); ftSZ or amplified tSZ autospectrum at 90 GHz (dotted blue curve), autospectrum at 150 GHz (dotted red curve), and 90 \times 150 GHz cross spectrum (dotted green curve). The combined CMB + ftSZ spectra at 90, 150, and 90 \times 150 GHz are also shown (solid blue, solid red, and solid green curves, respectively). Finally, the total power spectra at 90, 150, and 90 \times 150 GHz are shown (dashed cyan, dashed orange, and dashed lime curves, respectively). Of note is that there is no noise bias in the total power spectrum since we have used noise-debiased data-split cross spectra, though there are noise fluctuations.

amplitude parameters at the map level, but we choose to define them at the power spectrum level to be consistent with what is done in current major CMB data analyses (and also because the primary field of interest, the CMB, is entirely described by its power spectrum). We consider various methods for obtaining the posteriors on these parameters, outlined in the next subsection.

The tSZ field is highly non-Gaussian, as it traces hot gas in galaxy clusters. At the low multipole values considered in this work, the tSZ power spectrum lies significantly below that of the CMB. As described in Sec. I, our goal is to show that NILC decreases parameter error bars in the presence of non-Gaussian fields. To demonstrate that idea, we would need the tSZ power spectrum to be comparable to or larger than the CMB power spectrum in magnitude. Thus, each tSZ map is amplified by a factor of 150 for our simulations, since we are interested in showing that NILC-based summary statistics are an improvement over multifrequency power spectra summary statistics for data with large non-Gaussian components.

B. Outline of implemented methods

In this subsection we provide an outline of the methods considered in this work. We describe each method in detail in Secs. V–VII.

- (1) Multifrequency power spectra (see Sec. V)
 - (a) Gaussian likelihood (with known analytic parameter dependence)
 - (i) Maximum likelihood estimation (MLE)
 - (ii) Fisher matrix
 - (iii) Markov chain Monte Carlo (MCMC) algorithm
 - (b) Likelihood-free inference (LFI)
- (2) Harmonic ILC (HILC) power spectra (see Sec. VI)
 - (a) Gaussian likelihood
 - (i) Fixed weights
 - (A) Using known analytic parameter dependence
 - (I) MLE
 - (II) Fisher matrix
 - (III) MCMC algorithm
 - (B) Parameter dependence determined via symbolic regression (SR)
 - (I) MLE
 - (II) Fisher matrix
 - (III) MCMC algorithm
 - (ii) Varying weights
 - (A) Parameter dependence determined via SR
 - (I) MLE
 - (II) Fisher matrix
 - (III) MCMC algorithm
 - (b) LFI
 - (i) Fixed weights
 - (ii) Varying weights
- (3) Needlet ILC (NILC) power spectra (see Sec. VII)
 - (a) Gaussian likelihood (with parameter dependence determined via SR)
 - (i) MLE
 - (ii) Fisher matrix
 - (iii) MCMC algorithm
 - (b) LFI

V. MULTIFREQUENCY POWER SPECTRA AS SUMMARY STATISTICS

A. Gaussian likelihood

Fitting multifrequency autospectrum and cross-power spectrum data to parametric templates using a Gaussian likelihood is the classic method used in nearly all major primary CMB analyses to date. In our case, we use the autospectra and cross spectra at the two frequency channels as our data vector. Specifically, our data vector consists of $C_b^{90\text{split}1,90\text{split}2}$, $C_b^{90\text{split}1,150\text{split}2}$, $C_b^{150\text{split}1,90\text{split}2}$, and $C_b^{150\text{split}1,150\text{split}2}$, where 90split1 signifies the 90 GHz map with the first noise split, and similarly for the other frequency-split pairs. The subscript b denotes the multipole ℓ -space bin, where we use ten linearly spaced bins from $\ell = 2$ to $\ell = 250$ (one bin contains 24 multipoles, and all

other bins contain 25 multipoles). The Gaussian likelihood is given by

$$\ln \mathcal{L}(A_{\text{CMB}}, A_{\text{ftSZ}}) = -\frac{1}{2} \left(C_{b_1}^{ij}(A_{\text{CMB}}, A_{\text{ftSZ}}) - C_{b_1}^{ij,\text{data}} \right) \times (\text{Cov}_{ijb_1,klb_2})^{-1} \left(C_{b_2}^{kl}(A_{\text{CMB}}, A_{\text{ftSZ}}) - C_{b_2}^{kl,\text{data}} \right), \quad (16)$$

where there is implied summation over repeated indices i, j, k, l, b_1 , and b_2 on the rhs of Eq. (16). Here $C_{b_1}^{ij}, C_{b_2}^{kl} \in \{C_b^{90\text{split}1,90\text{split}2}, C_b^{90\text{split}1,150\text{split}2}, C_b^{150\text{split}1,90\text{split}2}, C_b^{150\text{split}1,150\text{split}2}\}$ and $\text{Cov}_{ijb_1,klb_2} \equiv \text{Cov}(C_{b_1}^{ij}, C_{b_2}^{kl})$ is the covariance matrix of the multifrequency power spectra, computed directly from 2000 independent simulations generated with the fiducial parameter values (i.e., there is no assumption of a Gaussian covariance matrix). We multiply the inverse of the multifrequency power spectrum covariance matrix by the correction factor $\frac{n-p-2}{n-1}$, where n is the number of simulations and p is the length of the data vector [66]. Since the number of simulations is far greater than the length of the data vector, this factor is very close to 1 in our case.

Note also that we have included off-diagonal ℓ bins in the covariance matrix and likelihood. For the multifrequency power spectra,

$$C_{b_1}^{ij}(A_{\text{CMB}}, A_{\text{ftSZ}}) = A_{\text{CMB}} C_{b_1}^{\text{CMB}} + A_{\text{ftSZ}} g_y^i g_y^j C_{b_1}^{\text{ftSZ}}, \quad (17)$$

where g_y^i and g_y^j are the spectral responses of the tSZ signal at frequencies i and j [as defined in Eq. (14)], $C_{b_1}^{\text{CMB}}$ is the fiducial CMB power spectrum template, and $C_{b_1}^{\text{ftSZ}}$ is the fiducial (amplified) Compton- y power spectrum template. For these fiducial templates, we use the mean of the CMB and ftSZ spectra generated from 2000 simulations to obtain smooth theory curves. We implement three methods for computing the posteriors on A_{CMB} and A_{ftSZ} : maximum likelihood estimation (MLE), the Fisher information matrix, and an MCMC algorithm.

For MLE, we numerically find the maximum-likelihood point for the likelihood function in Eq. (16) for each of the 2000 simulations. Specifically, for each simulation, we compute the power spectra of the frequency maps to use as $C_{b_1}^{ij,\text{data}}$ and $C_{b_2}^{kl,\text{data}}$. Note that these are noise debiased by construction, though there are noise fluctuations present. The total power spectrum at each frequency, along with the contributions from individual components, can be seen in Fig. 1. We then find the maximum-likelihood ($A_{\text{CMB}}, A_{\text{ftSZ}}$) point for every simulation using Nelder-Mead minimization [67] of the negative log likelihood. To avoid getting stuck in local minima, we set run the minimization routine using three starting points: $(A_{\text{CMB}}, A_{\text{ftSZ}}) = (1.0, 1.0)$, $(0.5, 0.5)$, and $(1.5, 1.5)$. We then use the global minimum

from these results as our maximum-likelihood point ($\hat{A}_{\text{CMB}}, \hat{A}_{\text{ftSZ}}$). The majority of the time, the results from each of these starting points are equal, and local minima do not pose a problem. Only in a few simulations are the results actually different depending on starting point.

For the Fisher matrix calculation, we obtain the parameter covariance matrix using the usual

$$\sigma_{\hat{A}_\alpha \hat{A}_\beta}^2 = (F_{\hat{A}_\alpha \hat{A}_\beta})^{-1}, \quad (18)$$

where $\hat{A}_\alpha, \hat{A}_\beta \in \{\hat{A}_{\text{CMB}}, \hat{A}_{\text{ftSZ}}\}$ and $F_{\hat{A}_\alpha \hat{A}_\beta}$ is the Fisher matrix defined as

$$\begin{aligned} F_{\hat{A}_\alpha \hat{A}_\beta} &= - \left(\frac{\partial^2 \ln \mathcal{L}}{\partial A_\alpha \partial A_\beta} \right) \Big|_{\hat{A}_\alpha \hat{A}_\beta} \\ &= \frac{\partial C_{b_1}^{ij}}{\partial A_\alpha} (\text{Cov}_{ijb_1, klb_2}^{-1}) \frac{\partial C_{b_2}^{kl}}{\partial A_\beta} \Big|_{A_\alpha = A_\beta = 1}, \end{aligned} \quad (19)$$

where there is implied summation over i, j, k, l, b_1 , and b_2 . As seen from Eq. (17), $\partial C_{b_1}^{ij} / \partial A_{\text{CMB}} = C_{b_1}^{\text{CMB}}$ and $\partial C_{b_1}^{ij} / \partial A_{\text{ftSZ}} = C_{b_1}^{\text{ftSZ}}$.

Finally, for the MCMC algorithm we use EMCEE [68] with the likelihood in Eq. (16) and an observation consisting of the mean of multifrequency autospectra and cross spectra over all simulations. We use ten walkers, each with a random starting point between $(A_{\text{CMB}}, A_{\text{ftSZ}}) = (0.8, 0.8)$ and $(A_{\text{CMB}}, A_{\text{ftSZ}}) = (1.2, 1.2)$. The autocorrelation time is used to assess MCMC convergence. Specifically, the chain is marked as converged when it is longer than 50 times the integrated autocorrelation time.

B. Likelihood-free inference

For LFI/SBI, we use the code SBI [69]. As discussed in Sec. III, we require four components for LFI: a prior, a simulator, an observation, and a network. For the former, we set a uniform prior centered on 1 with lower bound $(1 - 5\sigma_{\hat{A}_{\text{CMB}}}, 1 - 5\sigma_{\hat{A}_{\text{ftSZ}}})$ and upper bound $(1 + 5\sigma_{\hat{A}_{\text{CMB}}}, 1 + 5\sigma_{\hat{A}_{\text{ftSZ}}})$, where $\sigma_{\hat{A}_{\text{CMB}}}$ and $\sigma_{\hat{A}_{\text{ftSZ}}}$ are the 1D marginalized errors obtained from the multifrequency Gaussian likelihood on \hat{A}_{CMB} and \hat{A}_{ftSZ} , respectively. The 5σ range for the prior ensures that the prior is narrow enough to be efficient in the number of simulations, but wide enough to encompass the true error on the parameters (5σ of the results from the Gaussian likelihood should be sufficient since the Gaussian likelihood, though not entirely accurate, should not be a catastrophically bad approximation due to the central limit theorem). As a further check that the prior is wide enough, we ensure that the final posteriors do not hit the prior edges.

For the simulator, we draw A_{CMB} and A_{ftSZ} from the prior distribution. Then for each simulation, we multiply the input CMB map by a factor of $\sqrt{A_{\text{CMB}}}$ and the ftSZ map by

a factor of $\sqrt{A_{\text{ftSZ}}}$ (recall that the parameters are defined at the power spectrum level). We then proceed as usual with constructing maps of two noise splits for each frequency and constructing the data vector comprising the concatenation of $C_b^{90\text{split}1, 90\text{split}2}$, $C_b^{90\text{split}1, 150\text{split}2}$, $C_b^{150\text{split}1, 90\text{split}2}$, and $C_b^{150\text{split}1, 150\text{split}2}$. Thus, the data vector has a total length of 40 since there are ten bins in each of the concatenated spectra. We generate 40000 simulations with parameters drawn randomly from the prior. For the observation vector, we generate 2000 simulations of the data vector with $A_{\text{CMB}} = A_{\text{ftSZ}} = 1$ and take the mean over the 2000 simulations. This may result in a slightly biased central value of the posteriors since we use significantly fewer simulations to determine the mock observation vector than to determine the full posteriors. However, the focus of this paper is on the spread in the posteriors rather than central values. Note also that the observation vector here is allowed to be smooth since the simulator generates noisy spectra from which the LFI learns the covariance.

We use single-round neural posterior estimation (NPE) for the LFI. We apply z-score normalization to each simulated data vector and to the observation for more stable training. For the network, we train a masked autoregressive flow (MAF). We also experiment with using a neural spine flow, mixture density network, and masked autoencoder for distribution estimation, but we find that the MAF is the most efficient in the number of simulations for our purposes. We use 90% of the 40000 simulations for training and 10% for validation. We tune the hyperparameters of the network with the weights and biases framework (WANDB) [70]. Our WANDB project is public.³ We use a random hyperparameter search with 40 trials. For the final posteriors, we use an ensemble of the top ten networks in terms of logarithmic probability on the validation set and uniformly weight each of the top ten networks to obtain the final posteriors. To assess whether reductions in the size of the posterior are significant, we examine the variance in the final parameter covariance matrices obtained from different hyperparameter settings (the ‘‘error on the error’’). We find that this error on the error is not significant when using just the top ten networks. However, the final posteriors are still sensitive to the hyperparameter tuning, as some networks cause the posteriors to hit the prior edges and are thus essentially not learning any information. A table of hyperparameters that are tuned is shown in Table I. Specifically, we tune the learning rate, the number of transforms in the normalizing flow, the number of hidden features, the value at which to clip the total gradient norm in order to prevent exploding gradients, and the number of epochs to wait for improvement on the validation set before terminating training.

³https://wandb.ai/kmsurrao/cmb_sbi/sweeps.

TABLE I. Distributions of hyperparameters tuned in the MAF used for LFI. “Clip maximum norm” refers to the value at which to clip the total gradient norm in order to prevent exploding gradients, and “Epochs for convergence assessment” the number of epochs to wait for improvement on the validation set before terminating training. See Sec. VB for details of how the hyperparameter sweep is used. The same distributions of hyperparameters are used to tune the networks used for LFI with multifrequency power spectra, harmonic ILC power spectra, and needlet ILC power spectra summary statistics, described in Secs. VB, VIB, and VII B, respectively.

| Hyperparameter | Minimum value | Maximum value | Distribution |
|-----------------------------------|----------------------|----------------------|-----------------|
| Learning rate | 1.0×10^{-4} | 7.0×10^{-4} | Uniform |
| Number of transforms | 3 | 10 | Integer uniform |
| Hidden features | 35 | 65 | Integer uniform |
| Clip maximum norm | 3.0 | 7.0 | Uniform |
| Epochs for convergence assessment | 20 | 60 | Integer uniform |

VI. HARMONIC ILC COMPONENT MAP POWER SPECTRA AS SUMMARY STATISTICS

Our data vector for the harmonic ILC power spectra consists of $C_b^{\hat{T}\text{split1},\hat{T}\text{split2}}$, $C_b^{\hat{T}\text{split1},\hat{y}\text{split2}}$, $C_b^{\hat{y}\text{split1},\hat{T}\text{split2}}$, and $C_b^{\hat{y}\text{split1},\hat{y}\text{split2}}$, where $\hat{T}\text{split1}$ signifies a HILC CMB map produced using the frequency maps with the first noise split, $\hat{y}\text{split1}$ signifies a HILC Compton- γ map produced using the frequency maps with the first noise split, and similarly for the second split. The subscript b denotes the multipole ℓ -space bin, where we use ten linearly spaced bins from $\ell = 2$ to $\ell = 250$ (one bin contains 24 multipoles, and all other bins contain 25 multipoles).

We compute the HILC weights and power spectra analytically. As an example, to compute $C_b^{\hat{T}\text{split1},\hat{y}\text{split2}}$ we do the following. First, we produce four maps: a 90 GHz map with the first noise split, a 90 GHz map with the second noise split, a 150 GHz map with the first noise split, and a 150 GHz map with the second noise split. We then compute weights for the HILC CMB map using the first noise split in an analogous way to Eq. (3) [and noting the CMB spectral energy distribution (SED) is a vector of ones as described in Sec. IV]:

$$w_\ell^{i,\hat{T}\text{split1}} = \frac{\mathbb{1}_j(\hat{R}_\ell^{-1})_{ij}}{(\hat{R}_\ell^{-1})_{km} \mathbb{1}_k \mathbb{1}_m} \quad \text{with} \quad (\hat{R}_\ell)_{ij} = \sum_{\ell'=\ell-\Delta\ell/2}^{\ell+\Delta\ell/2} \frac{2\ell'+1}{4\pi} C_{\ell'}^{isplit1,jsplit1}, \quad (20)$$

where $isplit1, jsplit1 \in \{90split1, 150split1\}$, $\mathbb{1}$ is a vector of ones, and $\Delta\ell = 20$. This calculation is performed at every ℓ within windows that overlap, and then binned as described previously (note that the binning in the harmonic ILC computation at each ℓ is distinct from the binning of the final HILC spectra for use in the data vector). We then compute weights for the HILC Compton- γ map using the second noise split:

$$w_\ell^{i,\hat{y}\text{split2}} = \frac{(\hat{R}_\ell^{-1})_{ij} g_y^j}{(\hat{R}_\ell^{-1})_{km} g_y^k g_y^m} \quad \text{with} \quad (\hat{R}_\ell)_{ij} = \sum_{\ell'=\ell-\Delta\ell/2}^{\ell+\Delta\ell/2} \frac{2\ell'+1}{4\pi} C_{\ell'}^{isplit2,jsplit2}. \quad (21)$$

Then,

$$C_\ell^{\hat{T}\text{split1},\hat{y}\text{split2}} = \sum_{i,j} w_\ell^{i,\hat{T}\text{split1}} w_\ell^{j,\hat{y}\text{split2}} C_\ell^{isplit1,jsplit2}. \quad (22)$$

After computing the spectrum independently for each multipole, ℓ -space binning is performed, giving $C_b^{\hat{T}\text{split1},\hat{y}\text{split2}}$. The computation of $C_b^{\hat{T}\text{split1},\hat{T}\text{split2}}$, $C_b^{\hat{y}\text{split1},\hat{T}\text{split2}}$, and $C_b^{\hat{y}\text{split1},\hat{y}\text{split2}}$ are done analogously. In Appendix A we verify that this analytic computation of HILC spectra yields the same results as producing HILC maps and then computing the spectra from the maps.

We use 2000 simulations for the HILC Gaussian likelihood pipelines (described in Sec. VIA). We have two options for computing the weights in the pipeline:

- (1) *Fixed weights.* Compute the HILC weights once from some fiducial template multifrequency power spectra. Apply these same weights to every simulation. For the remainder of the paper, we refer to this as the “fixed weights” or “weights once” HILC variant.
- (2) *Varying weights.* Compute the HILC weights separately for every simulation, using the specific multifrequency power spectra from that simulation. For the remainder of the paper, we refer to this as the “varying weights” or “weights vary” HILC variant.

When the weights vary on a per-simulation basis, we ensure that the ILC bias is mitigated by omitting the central ℓ value in each bin when computing the weights, i.e., by setting

$$(\hat{R}_\ell)_{ij} = \left(\sum_{\ell'=\ell-\Delta\ell/2}^{\ell+\Delta\ell/2} \frac{2\ell'+1}{4\pi} C_\ell^{ij} \right) - \frac{2\ell+1}{4\pi} C_\ell^{ij}. \quad (23)$$

This heavily reduces chance correlations between the signal and noise/contaminants, which thereby reduces the ILC bias (this is analogous to a similar trick developed for NILC in Ref. [16]). The two variants of the HILC pipeline have different implications for parameter dependence in the Gaussian likelihood. This is discussed further below in Sec. VI A.

A. Gaussian likelihood

The Gaussian likelihood is given by

$$\begin{aligned} \ln \mathcal{L}(A_{\text{CMB}}, A_{\text{fitSZ}}) &= -\frac{1}{2} (C_{b_1}^{\hat{p}\hat{q}}(A_{\text{CMB}}, A_{\text{fitSZ}}) - C_{b_1}^{\hat{p}\hat{q}, \text{data}}) \\ &\quad \times (\text{Cov}_{\hat{p}\hat{q}b_1, \hat{r}\hat{s}b_2})^{-1} (C_{b_2}^{\hat{r}\hat{s}}(A_{\text{CMB}}, A_{\text{fitSZ}}) \\ &\quad - C_{b_2}^{\hat{r}\hat{s}, \text{data}}), \end{aligned} \quad (24)$$

where summation is implied over repeated indices \hat{p} , \hat{q} , \hat{r} , \hat{s} , b_1 , and b_2 on the rhs. Here $C_{b_1}^{\hat{p}\hat{q}}, C_{b_2}^{\hat{r}\hat{s}} \in \{C_b^{\hat{T}\text{split}1, \hat{T}\text{split}2}, C_b^{\hat{T}\text{split}1, \hat{y}\text{split}2}, C_b^{\hat{y}\text{split}1, \hat{T}\text{split}2}, C_b^{\hat{y}\text{split}1, \hat{y}\text{split}2}\}$ and $\text{Cov}_{\hat{p}\hat{q}b_1, \hat{r}\hat{s}b_2} \equiv \text{Cov}(C_{b_1}^{\hat{p}\hat{q}}, C_{b_2}^{\hat{r}\hat{s}})$ is the covariance matrix of the HILC power spectra, computed directly from 2000 independent simulations generated at the fiducial parameter values (i.e., there is no assumption of a Gaussian covariance matrix). A correction factor is applied to the inverse of the HILC power spectrum covariance matrix, as is done in Sec. VA for the multifrequency power spectrum covariance matrix [66]. Note also that we have included off-diagonal ℓ bins in the covariance matrix and likelihood.

1. Analytic parameter dependence

When using fixed weights in the HILC pipeline, the weights do not depend on the particular fluctuations of any single simulated sky realization. Thus, we can analytically write down the parameter dependence as (ignoring binning for now)

$$\begin{aligned} C_\ell^{\hat{p}\hat{q}} &= \sum_{i,j} w_\ell^{i,\hat{p}} w_\ell^{j,\hat{q}} C_\ell^{ij}(A_{\text{CMB}}, A_{\text{fitSZ}}) \\ &= \sum_{i,j} w_\ell^{i,\hat{p}} w_\ell^{j,\hat{q}} (A_{\text{CMB}} C_\ell^{\text{CMB}} + A_{\text{fitSZ}} g_y^i g_y^j C_\ell^{\text{fitSZ}}), \end{aligned} \quad (25)$$

where we have used Eq. (17) in the second equality. Including binning, we obtain

$$\begin{aligned} C_{b_1}^{\hat{p}\hat{q}} &= \frac{1}{N_{b_1}} \sum_{\ell \in b_1} \sum_{i,j} w_\ell^{i,\hat{p}} w_\ell^{j,\hat{q}} C_\ell^{ij}(A_{\text{CMB}}, A_{\text{fitSZ}}) \\ &= \frac{1}{N_{b_1}} \sum_{\ell \in b_1} \sum_{i,j} w_\ell^{i,\hat{p}} w_\ell^{j,\hat{q}} (A_{\text{CMB}} C_\ell^{\text{CMB}} + A_{\text{fitSZ}} g_y^i g_y^j C_\ell^{\text{fitSZ}}), \end{aligned} \quad (26)$$

where N_{b_1} is the number of multipole values ℓ contained in bin b_1 [it is *not* the total number of modes (ℓ, m)].

However, this simple parameter dependence is not correct in the varying-weights case. In that case, the weights respond to the particular fluctuations in the sky signal realized in each simulation. Thus, they are functions of the (randomly fluctuating) parameter amplitude values of individual simulations. Determining the parameter dependence in the varying-weights case is the subject of Sec. VI A 2.

2. Parameter dependence from symbolic regression

Our goal in this section is to write

$$\begin{aligned} C_b^{\hat{p}\hat{q}}(A_{\text{CMB}}, A_{\text{fitSZ}}) &= f_b^{\hat{p}\hat{q}}(A_{\text{CMB}}, A_{\text{fitSZ}}) \\ &\quad \times C_b^{\hat{p}\hat{q}}(A_{\text{CMB}} = 1 = A_{\text{fitSZ}}) \end{aligned} \quad (27)$$

and determine a function $f_b^{\hat{p}\hat{q}}(A_{\text{CMB}}, A_{\text{fitSZ}})$ for each $\hat{p}\hat{q} \in \{\hat{T}\text{split}1\hat{T}\text{split}2, \hat{T}\text{split}1\hat{y}\text{split}2, \hat{y}\text{split}1\hat{T}\text{split}2, \hat{y}\text{split}1\hat{y}\text{split}2\}$ and bin b . The superscripts and subscripts on f only denote that it is a quantity to be determined separately for each $\hat{p}\hat{q}$ and b ; however, f is just some function, not a power spectrum. Moreover, there is no implied summation over \hat{p} , \hat{q} , or b in the equation. To avoid assuming a specified form of f , we find these functions with symbolic regression, using PySR [71]. For our regressor, we allow the binary operations of addition, subtraction, multiplication, and division; unary operations of exponentiation, squaring, and cubing; and a custom inverse operation. We use an L2-loss function that does not in itself penalize complexity since we simply need smooth functions that describe the simulations well, rather than concise equations representing any fundamental law of physics. However, we do impose a complexity limit to avoid overfitting by setting the maximum size of our expression to 12.

For each $\hat{p}\hat{q}$ and bin b , the values that we feed into the regressor come from scaling various combinations of components. We define four ‘‘scaling factors’’ around the fiducial parameter value of 1: 0.9, 0.99, 1.01, and 1.1. First, we find the average HILC spectra over 50 simulations and use these to compute

$$C_b^{\hat{p}\hat{q}}(A_{\text{CMB}} = 1 = A_{\text{fitSZ}}).$$

As an example, consider the scaling factor 0.9. Using the same map-level realizations that were used to determine

$C_b^{\hat{p}\hat{q}}(A_{\text{CMB}} = 1 = A_{\text{fSZ}})$, we multiply the CMB and/or the fSZ map by 0.9 for each of the 50 simulations. The ILC weights (in the varying weights case) also change as a result. Then we compute

$$C_b^{\hat{p}\hat{q}}(A_{\text{CMB}} = 0.81, A_{\text{fSZ}} = 1), \quad C_b^{\hat{p}\hat{q}}(A_{\text{CMB}} = 1, A_{\text{fSZ}} = 0.81), \quad \text{and} \quad C_b^{\hat{p}\hat{q}}(A_{\text{CMB}} = 0.81, A_{\text{fSZ}} = 0.81).$$

Note that the scaled parameters are set to 0.81 since multiplying by 0.9 at the map level corresponds to multiplying by $0.9^2 = 0.81$ at the power spectrum level. For this example, the data points (\vec{x}, y) that go into the regressor are then

$$\begin{aligned} \vec{x} &= (A_{\text{CMB}}, A_{\text{fSZ}}) = (0.81, 1), & y &= C_b^{\hat{p}\hat{q}}(A_{\text{CMB}} = 0.81, A_{\text{fSZ}} = 1) / C_b^{\hat{p}\hat{q}}(A_{\text{CMB}} = 1 = A_{\text{fSZ}}), \\ \vec{x} &= (A_{\text{CMB}}, A_{\text{fSZ}}) = (1, 0.81), & y &= C_b^{\hat{p}\hat{q}}(A_{\text{CMB}} = 1, A_{\text{fSZ}} = 0.81) / C_b^{\hat{p}\hat{q}}(A_{\text{CMB}} = 1 = A_{\text{fSZ}}), \quad \text{and} \\ \vec{x} &= (A_{\text{CMB}}, A_{\text{fSZ}}) = (0.81, 0.81), & y &= C_b^{\hat{p}\hat{q}}(A_{\text{CMB}} = 0.81 = A_{\text{fSZ}}) / C_b^{\hat{p}\hat{q}}(A_{\text{CMB}} = 1 = A_{\text{fSZ}}). \end{aligned}$$

We repeat the same procedure with the other scaling factors, feeding the (\vec{x}, y) pairs from all scaling factors into the regressor together. Then by definition of \vec{x} and y , the regressor learns $f_b^{\hat{p}\hat{q}}$.

This procedure is used to determine the parameter dependence for the HILC varying-weights Gaussian-likelihood pipeline. However, it can also be used on the HILC fixed-weights Gaussian-likelihood pipeline. Since we know the analytic parameter dependence in the fixed-weights case, for that scenario we can compare the posteriors obtained using analytic parameter dependence to the posteriors obtained using SR as a cross-check of the validity of the SR approach. This serves as an important validation step for the varying-weights HILC Gaussian-likelihood pipeline, as well as the NILC Gaussian-likelihood pipeline that is described in Sec. VII A. Sample equations and parameter dependence learned via SR are shown in Appendix B. Note that, because of the way we define the parameter dependence function $f_b^{\hat{p}\hat{q}}$ in Eq. (27), we cannot directly compare the equations obtained from symbolic regression to the analytic parameter dependence in Eq. (26). We could have instead defined the symbolic regression procedure such that it could recover that same analytic parameter dependence, but the computational cost would be greater since we would have to perform symbolic regression multiple times for each \hat{p} , \hat{q} , and b .

3. MLE, Fisher matrix, and MCMC

As for the multifrequency power spectrum Gaussian likelihood, we obtain parameter posteriors from HILC map power spectra via MLE, the Fisher matrix, and an MCMC algorithm. For MLE, we numerically find the maximum-likelihood point of the likelihood function in Eq. (24) for each of the 2000 simulations. Specifically, for each simulation, we compute the HILC spectra from that simulation to use as $C_{b_1}^{\hat{p}\hat{q}, \text{data}}$ and $C_{b_2}^{\hat{r}\hat{s}, \text{data}}$. We then find the maximum-likelihood $(A_{\text{CMB}}, A_{\text{fSZ}})$ point for every simulation using Nelder-Mead minimization [67] of the

negative log likelihood. We run the minimization routine using three starting points: $(A_{\text{CMB}}, A_{\text{fSZ}}) = (1.0, 1.0)$, $(0.5, 0.5)$, and $(1.5, 1.5)$. We then use the global minimum from these results as our maximum-likelihood point $(\hat{A}_{\text{CMB}}, \hat{A}_{\text{fSZ}})$.

The Fisher matrix in this approach is defined as

$$\begin{aligned} F_{\hat{A}_\alpha \hat{A}_\beta} &= - \left(\frac{\partial^2 \ln \mathcal{L}}{\partial A_\alpha \partial A_\beta} \right) \Big|_{\hat{A}_\alpha \hat{A}_\beta} \\ &= \frac{\partial C_{b_1}^{\hat{p}\hat{q}}}{\partial A_\alpha} (\text{Cov}_{\hat{p}\hat{q}b_1, \hat{r}\hat{s}b_2}^{-1}) \frac{\partial C_{b_2}^{\hat{r}\hat{s}}}{\partial A_\beta} \Big|_{A_\alpha = A_\beta = 1}, \quad (28) \end{aligned}$$

where there is implied summation over \hat{p} , \hat{q} , \hat{r} , \hat{s} , b_1 , and b_2 . For the case of fixed weights with analytic parameter dependence, we compute the derivatives in Eq. (28) analytically [cf. Eq. (26)]. For varying weights and fixed weights with parameter dependence determined via SR, we compute the derivatives numerically with finite differences.

Finally, for the MCMC algorithm we use EMCEE [68] with the likelihood in Eq. (24) and a mock observation consisting of the mean of simulated HILC autospectra and cross spectra over all simulations. The remaining setup of the MCMC algorithm is the same as described in Sec. VA.

B. Likelihood-free inference

We follow a similar procedure as for the LFI in the multifrequency power spectrum case, described in Sec. VB. For the prior, we set a uniform prior centered on 1 with lower bound $(1 - 5\sigma_{\hat{A}_{\text{CMB}}}, 1 - 5\sigma_{\hat{A}_{\text{fSZ}}})$ and upper bound $(1 + 5\sigma_{\hat{A}_{\text{CMB}}}, 1 + 5\sigma_{\hat{A}_{\text{fSZ}}})$, where $\sigma_{\hat{A}_{\text{CMB}}}$ and $\sigma_{\hat{A}_{\text{fSZ}}}$ are the 1D marginalized errors obtained from the multifrequency Gaussian likelihood on \hat{A}_{CMB} and \hat{A}_{fSZ} , respectively. Note that the prior is kept the same as the prior in the multifrequency LFI for fair comparison of the resulting posteriors.

As for the HILC Gaussian likelihood, we perform LFI with two versions of the HILC method: one in which the weights are computed once from some fiducial template

generated at $A_{\text{CMB}} = 1 = A_{\text{ftSZ}}$, and one in which the weights are computed independently for each output of the simulator. For the simulator, we draw A_{CMB} and A_{ftSZ} from the prior distribution. Then for each simulation, we multiply the input CMB map by a factor of $\sqrt{A_{\text{CMB}}}$ and the ftSZ map by a factor of $\sqrt{A_{\text{ftSZ}}}$. We then proceed as usual with creating maps of two noise splits for each frequency and performing the HILC operations to construct the data vector comprising the concatenation of $C_b^{\hat{T}\text{split1},\hat{T}\text{split2}}$, $C_b^{\hat{T}\text{split1},\hat{y}\text{split2}}$, $C_b^{\hat{y}\text{split1},\hat{T}\text{split2}}$, and $C_b^{\hat{y}\text{split1},\hat{y}\text{split2}}$. Thus, the data vector has a total length of 40 since there are ten bins in each of the concatenated spectra. The remainder of the LFI setup is the same as described in Sec. VB.

VII. NEEDLET ILC COMPONENT MAP POWER SPECTRA AS SUMMARY STATISTICS

A. Gaussian likelihood

We follow a procedure similar to that described in Sec. VIA. Our data vector for the NILC component map autospectra and cross-power spectra consists of $C_b^{\hat{T}\text{split1},\hat{T}\text{split2}}$, $C_b^{\hat{T}\text{split1},\hat{y}\text{split2}}$, $C_b^{\hat{y}\text{split1},\hat{T}\text{split2}}$, and $C_b^{\hat{y}\text{split1},\hat{y}\text{split2}}$, where $\hat{T}\text{split1}$ signifies a NILC CMB map produced using the frequency maps with the first noise split, $\hat{y}\text{split1}$ signifies a NILC Compton- y map produced using the frequency maps with the first noise split, and similarly for the second split. The subscript b denotes the multipole bin, where we use ten bins from $\ell = 2$ to $\ell = 250$ (one bin contains 24 multipoles, and all other bins contain 25 multipoles).

In the HILC case, we are able to easily compute the power spectra analytically. However, this is not the case for NILC, as shown in Ref. [24], due to the nontrivial contributions of three- and four-point functions of the components and weight maps. Thus, we build NILC maps from simulations and compute the power spectra directly from these. To produce the NILC maps and weight maps, we use PYILC⁴ [17,36]. We use needlet filters that are the difference of successive Gaussians, where the Gaussians have full width at half maximum (FWHM) values of 300, 120, and 60 arcmin, resulting in four needlet filter scales. The fourth filter is modified slightly such that the filters satisfy the needlet filter power preservation constraint, i.e., that $\sum_{(n)} (h_\ell^{(n)})^2 = 1$ at each ℓ . The filters are shown in Fig. 2. Examples of the resulting NILC component maps and weight maps are shown in Appendix A.

In the HILC case, we implement two variants: one with fixed ILC weights (i.e., weights that are not recomputed for every simulation realization) and one with varying ILC weights (i.e., weights that are recomputed for every realization). In the NILC case, we only use varying




FIG. 2. Needlet filters $h_\ell^{(n)}$ used in our simulations for the NILC pipeline, found by taking the differences of successive Gaussians of FWHM 300, 120, and 60 arcmin.

weights, meaning that the weight maps are determined separately for every simulation. This is necessary so that the NILC algorithm can capture meaningful spatially varying information within each simulation. In particular, on a single simulation, the pipeline is as follows:

- (1) Simulate four frequency maps (maps of two noise splits for each of the two frequencies).
- (2) For each split, run PYILC to build a NILC CMB-preserved map (e.g., to build a CMB NILC map for the first split, feed in the 90 GHz Split 1 and 150 GHz Split 1 maps as input to the code and specify preservation of the CMB component). Repeat the same procedure for the (amplified) Compton- y NILC map. Thus, for each simulation we have four NILC maps: $\hat{T}\text{split1}$, $\hat{T}\text{split2}$, $\hat{y}\text{split1}$, and $\hat{y}\text{split2}$.
- (3) Compute the elements of the data vector (i.e., the power spectra): $C_b^{\hat{T}\text{split1},\hat{T}\text{split2}}$, $C_b^{\hat{T}\text{split1},\hat{y}\text{split2}}$, $C_b^{\hat{y}\text{split1},\hat{T}\text{split2}}$, and $C_b^{\hat{y}\text{split1},\hat{y}\text{split2}}$.

We then use the same Gaussian likelihood as in Eq. (24), except that the spectra are now NILC spectra instead of HILC spectra. To determine the parameter dependence $C_b^{\hat{p}\hat{q}}(A_{\text{CMB}}, A_{\text{ftSZ}})$, we follow the same symbolic regression procedure that is described in Sec. VIA 2. In Appendix B, we show examples of the expressions and parameter dependence we obtain from symbolic regression, and compare them to those obtained in the HILC case.

We compute the covariance matrix used in the Gaussian likelihood directly via 2000 simulations of the above pipeline. A correction factor is applied to the inverse of the NILC power spectrum covariance matrix, as is done in Sec. VA for the multifrequency power spectrum covariance matrix [66]. We then proceed with the MLE, Fisher matrix, and MCMC procedures, as described in Sec. VIA 3. In Fig. 3, we show the correlation matrix and covariance matrix from these NILC simulations and compare them to the HILC correlation and covariance matrices for both the fixed-weights and varying-weights cases

⁴<https://github.com/jcolinhill/pyilc>.




FIG. 3. Correlation matrices (top row) and absolute value of covariance matrices (bottom row) computed via 2000 simulations, as used in our Gaussian likelihoods. These matrices are shown for the NILC pipeline (left), HILC with varying weights pipeline (middle), and HILC with fixed weights pipeline (right). Within each subplot, there is a grid of nine squares, corresponding to elements of the covariance and correlation among $C_{\ell}^{\hat{T}\hat{T}}$, $C_{\ell}^{\hat{T}\hat{y}}$, and $C_{\ell}^{\hat{y}\hat{y}}$ for each of the ten multipole bins increasing from left to right and top to bottom. Here $C_{\ell}^{\hat{T}\hat{y}}$ corresponds to $C_{\ell}^{\hat{T}\text{split1},\hat{y}\text{split2}}$ and we omit the other split combination for concision since it behaves similarly. The \hat{T} maps are in units of K_{CMB} , while the \hat{y} maps are in dimensionless Compton-y units.

(described in Sec. VI). In all cases, the non-Gaussianity of the tSZ field is evident from the off-diagonal contributions to $\text{Cov}(C_{\ell}^{\hat{y}\hat{y}}, C_{\ell}^{\hat{y}\hat{y}})$, and the propagation of that field as a contaminant to the $\text{Cov}(C_{\ell}^{\hat{T}\hat{T}}, C_{\ell}^{\hat{T}\hat{T}})$ covariance can be seen. However, we see that this contamination is lowest in the NILC case, where the $\text{Cov}(C_{\ell}^{\hat{T}\hat{T}}, C_{\ell}^{\hat{T}\hat{T}})$ is roughly diagonal. We also note that the varying-weights HILC does better than the fixed-weights HILC, as expected since the varying weights can account for (isotropic) harmonic-space fluctuations in individual simulations. As further confirmation that NILC better cleans contaminants, we show the mean HILC (in the varying-weights case) and mean NILC spectra from 2000 simulations in Fig. 4. From this figure, it is evident that NILC better suppresses contamination and recovers spectra that are closer to the underlying ground truth.

B. Likelihood-free inference

We follow a similar procedure to that used for the LFI in the multifrequency-power-spectra and HILC-power-spectra

cases, described in Secs. VB and VI B, respectively. For the prior, we set a uniform prior centered on 1 with lower bound $(1 - 5\sigma_{\hat{A}_{\text{CMB}}}, 1 - 5\sigma_{\hat{A}_{\text{tSZ}}})$ and upper bound $(1 + 5\sigma_{\hat{A}_{\text{CMB}}}, 1 + 5\sigma_{\hat{A}_{\text{tSZ}}})$, where $\sigma_{\hat{A}_{\text{CMB}}}$ and $\sigma_{\hat{A}_{\text{tSZ}}}$ are the 1D marginalized errors obtained from the multifrequency Gaussian likelihood on \hat{A}_{CMB} and \hat{A}_{tSZ} , respectively. Note that the prior is kept the same as the prior in the multifrequency and HILC LFI for fair comparison of the resulting posteriors.

For the simulator, we draw A_{CMB} and A_{tSZ} from the prior distribution. Then for each simulation, we multiply the input CMB map by a factor of $\sqrt{A_{\text{CMB}}}$ and the ftSZ map by a factor of $\sqrt{A_{\text{tSZ}}}$. We then proceed as in the Gaussian-likelihood case with creating maps of two noise splits for each frequency and performing the NILC operations to construct the data vector comprising the concatenation of $C_b^{\hat{T}\text{split1},\hat{T}\text{split2}}$, $C_b^{\hat{T}\text{split1},\hat{y}\text{split2}}$, $C_b^{\hat{y}\text{split1},\hat{T}\text{split2}}$, and $C_b^{\hat{y}\text{split1},\hat{y}\text{split2}}$. Thus, the data vector has a total length of 40 since there are ten bins in each of the concatenated spectra. The remainder of the LFI setup is the same as described in Sec. VB.




FIG. 4. Mean ILC spectra over 2000 simulations, for NILC and HILC (with varying weights, i.e., weights determined separately for each of the simulations). Solid curves show the ground truth spectra for the CMB (blue), CMB \times amplified Compton-y (red), and amplified Compton-y (green). Dashed curves show HILC spectra: $C_\ell^{\hat{T}\hat{T}}$ (light blue), $C_\ell^{\hat{T}\hat{y}}$ (orange), and $C_\ell^{\hat{y}\hat{y}}$ (lime green). Dotted curves show NILC spectra: $C_\ell^{\hat{T}\hat{T}}$ (cyan), $C_\ell^{\hat{T}\hat{y}}$ (salmon), and $C_\ell^{\hat{y}\hat{y}}$ (olive). Here \hat{T} represents a CMB-preserved ILC map, and \hat{y} represents an ILC map that preserves the tSZ signal. Each of the spectra are computed using noise splits (here $C_\ell^{\hat{T}\hat{T}}$ is really $C_\ell^{\hat{T}\text{split1},\hat{y}\text{split2}}$ and the other split combination is omitted for concision since it behaves similarly). The \hat{T} maps are in units of K_{CMB} , and the \hat{y} maps are in dimensionless Compton-y units.

VIII. WHEN ARE THE HARMONIC ILC AND MULTIFREQUENCY POWER SPECTRUM INFERENCE METHODS ANALYTICALLY EQUIVALENT?

A. General setup

In this section, we perform analytic calculations to compare parameter covariances derived from HILC map power spectra and template fitting to multifrequency autospectra and cross-power spectra. We determine the conditions under which these two inference methods are exactly equivalent, and explicitly prove this result. For all calculations here, we assume a Gaussian likelihood for both approaches. We also assume no ℓ -space binning, but the results also hold in the limit of small bin widths. Finally, for clarity and concision we do not use different noise splits of the data, though the results hold in that case as well.

We prove the equivalence of the parameter Fisher matrices, and thus the final parameter covariance matrices, at a single multipole ℓ (or for a single narrow bandpower). This immediately generalizes to fitting all multipoles simultaneously, as long as the power spectrum covariance matrix is diagonal in ℓ . The Fisher matrix for the multifrequency approach is given by Eq. (19):

$$F_{\hat{A}_\alpha \hat{A}_\beta}^{\text{mult}} = \sum_{ij} \sum_{km} \frac{\partial C_\ell^{ij}}{\partial A_\alpha} (\text{Cov}^{-1})_{\ell}^{ij,km} \frac{\partial C_\ell^{km}}{\partial A_\beta} \Big|_{A_\alpha=1=A_\beta}, \quad (29)$$

where ij, km can each take on $N_{\text{freq}}(N_{\text{freq}} + 1)/2$ values. For example, if there are two frequency channels, then ij, km can each take on three values: one for each frequency-frequency autospectrum, and one for the frequency-frequency cross-spectrum. Since we consider only a single ℓ here, there is no sum over ℓ, ℓ' , as would be present in the general form of the above equation. The Fisher matrix for the HILC approach is given by Eq. (28):

$$F_{\hat{A}_\alpha \hat{A}_\beta}^{\text{HILC}} = \sum_{\hat{p}\hat{q}} \sum_{\hat{r}\hat{s}} \frac{\partial C_\ell^{\hat{p}\hat{q}}}{\partial A_\alpha} (\text{Cov}^{-1})_{\ell}^{\hat{p}\hat{q},\hat{r}\hat{s}} \frac{\partial C_\ell^{\hat{r}\hat{s}}}{\partial A_\beta} \Big|_{A_\alpha=1=A_\beta}, \quad (30)$$

where the number of values that $\hat{p}\hat{q}, \hat{r}\hat{s}$ can take on is more complicated than in the multifrequency case and is examined in detail in the next subsection. The power spectra of HILC maps are easy to model. In particular, from Eq. (25), we have

$$\frac{\partial C_\ell^{\hat{p}\hat{q}}}{\partial A_\alpha} = \sum_{i,j} w_\ell^{i,\hat{p}} w_\ell^{j,\hat{q}} \frac{\partial C_\ell^{ij}}{\partial A_\alpha} \quad \text{and} \quad \frac{\partial C_\ell^{\hat{r}\hat{s}}}{\partial A_\beta} = \sum_{k,l} w_\ell^{k,\hat{r}} w_\ell^{l,\hat{s}} \frac{\partial C_\ell^{kl}}{\partial A_\beta}. \quad (31)$$

The covariance matrix for the HILC component map autospectra and cross-power spectra is (at a single ℓ)

$$\begin{aligned} \text{Cov}_{\ell}^{\hat{p}\hat{q},\hat{r}\hat{s}} &= \text{Cov}(C_{\ell}^{\hat{p}\hat{q}}, C_{\ell}^{\hat{r}\hat{s}}) \\ &= \sum_{i,j,k,m} w_{\ell}^{i,\hat{p}} w_{\ell}^{j,\hat{q}} w_{\ell}^{k,\hat{r}} w_{\ell}^{m,\hat{s}} \text{Cov}(C_{\ell}^{ij}, C_{\ell}^{km}), \end{aligned} \quad (32)$$

where here we are summing over individual frequencies, i.e., i, j, k, l each take on N_{freq} possible values. Then we have

$$\begin{aligned} F_{\hat{A}_\alpha \hat{A}_\beta}^{\text{HILC}} &= \sum_{\hat{p}\hat{q},\hat{r}\hat{s}} \left(\sum_{i,j} w_{\ell}^{i,\hat{p}} w_{\ell}^{j,\hat{q}} \frac{\partial C_{\ell}^{ij}}{\partial A_{\alpha}} \right) \\ &\times \left(\sum_{i,j,k,l} w_{\ell}^{i,\hat{p}} w_{\ell}^{j,\hat{q}} w_{\ell}^{k,\hat{r}} w_{\ell}^{l,\hat{s}} \text{Cov}(C_{\ell}^{ij}, C_{\ell}^{kl}) \right)^{-1} \\ &\times \left(\sum_{k,l} w_{\ell}^{k,\hat{r}} w_{\ell}^{l,\hat{s}} \frac{\partial C_{\ell}^{kl}}{\partial A_{\beta}} \right) \Big|_{A_{\alpha}=1=A_{\beta}}. \end{aligned} \quad (33)$$

Here we have assumed fixed HILC weights, e.g., as would be determined once from the actual data and applied to every simulation. Thus, the weights are not parameter-dependent quantities [if they were, then the derivatives in Eq. (31) would be more complicated]. Note that there are some subtleties in the indexing. In particular, we must define some mapping from $\hat{p}\hat{q}$ to individual \hat{p} and \hat{q} , and similarly for $\hat{r}\hat{s}$. This is discussed in the next subsection.

B. Number of components vs number of frequencies

We prove the equivalence of $F_{\hat{A}_\alpha \hat{A}_\beta}^{\text{mult}}$ and $F_{\hat{A}_\alpha \hat{A}_\beta}^{\text{HILC}}$ analytically for specific numbers of frequencies and sky components. In all cases, we assume an arbitrary frequency-frequency covariance matrix that is diagonal in ℓ , unless otherwise noted. The Gaussian covariance matrix as in Refs. [72,73] is an example of such a covariance matrix. Because of the lengthy algebraic expressions in the following analysis, we provide *Mathematica* notebooks demonstrating each of these results in our GitHub repository NILC-Inference-Pipeline.

1. Two components, two frequencies

This is the simplest case and is also the one we consider in our full simulation-based pipelines. We assume an arbitrary diagonal frequency-frequency power spectrum covariance matrix for this proof, but we do not assume that this covariance matrix is Gaussian. For concreteness, let us assume two frequency channels, e.g., 90 and 150 GHz. Moreover, we assume two sky components, the CMB and tSZ fields, and their respective minimum-variance HILC maps, \hat{T} and \hat{y} . In this case, we have that $ij, km \in \{90 \times 90, 90 \times 150, 150 \times 150\}$. This yields a rank-3 frequency-frequency covariance matrix at each ℓ , $\text{Cov}_{\ell}^{ij,km}$. Since there are two components, we have $\hat{p}\hat{q} \in \{\hat{T}\hat{T}, \hat{T}\hat{y}, \hat{y}\hat{y}\}$. The covariance matrix $\text{Cov}_{\ell}^{\hat{p}\hat{q},\hat{r}\hat{s}}$ has a rank of 3, as can immediately be seen via the construction in Eq. (32) in terms of $\text{Cov}_{\ell}^{ij,km}$. Thus, the HILC data vector

comprises all three spectra $C_{\ell}^{\hat{T}\hat{T}}, C_{\ell}^{\hat{T}\hat{y}}, C_{\ell}^{\hat{y}\hat{y}}$. Using that HILC data vector, the parameter covariance matrix elements from HILC and multifrequency power spectra are equal.

2. Three components, two frequencies

Let us now suppose that we have three components whose HILC maps are denoted \hat{T}, \hat{y} , and \hat{J} . Suppose we let $C_{\ell}^{\hat{p}\hat{q}}$ take on values of all the component autospectra and cross spectra, i.e., $C_{\ell}^{\hat{p}\hat{q}} \in \{C_{\ell}^{\hat{T}\hat{T}}, C_{\ell}^{\hat{T}\hat{y}}, C_{\ell}^{\hat{T}\hat{J}}, C_{\ell}^{\hat{y}\hat{y}}, C_{\ell}^{\hat{y}\hat{J}}, C_{\ell}^{\hat{J}\hat{J}}\}$. We would then have a 6×6 HILC power spectrum covariance matrix. However, this matrix actually only has a rank of 3 since it is built from the rank-3 frequency-frequency power spectrum covariance matrix, and we must thus remove linearly dependent rows and columns to make the covariance matrix invertible. In particular, the rows and columns corresponding to the elements of the data vector in the following sets are linearly dependent: $\{C_{\ell}^{\hat{T}\hat{T}}, C_{\ell}^{\hat{T}\hat{y}}, C_{\ell}^{\hat{T}\hat{J}}\}$, $\{C_{\ell}^{\hat{y}\hat{y}}, C_{\ell}^{\hat{y}\hat{J}}, C_{\ell}^{\hat{J}\hat{J}}\}$, and $\{C_{\ell}^{\hat{J}\hat{J}}, C_{\ell}^{\hat{T}\hat{J}}, C_{\ell}^{\hat{y}\hat{J}}\}$. Using $C_{\ell}^{\hat{p}\hat{q}} \in \{C_{\ell}^{\hat{T}\hat{T}}, C_{\ell}^{\hat{y}\hat{y}}, C_{\ell}^{\hat{J}\hat{J}}\}$ as our fundamental HILC data vector thus solves this problem. Using this construction in Eq. (33), we find that $F_{\hat{A}_\alpha \hat{A}_\beta}^{\text{HILC}} = F_{\hat{A}_\alpha \hat{A}_\beta}^{\text{mult}}$, and thus the HILC and multifrequency-power-spectrum methods yield identical parameter covariance matrix elements. (The analytic demonstration can be found in the *Mathematica* notebooks in NILC-Inference-Pipeline.)

3. Two components, three frequencies

For concreteness, suppose that the three frequency channels are 90, 150, and 220 GHz. Then there are six possible values for each of ij and km : $ij, km \in \{90 \times 90, 150 \times 150, 220 \times 220, 90 \times 150, 90 \times 220, 150 \times 220\}$. The frequency-frequency power spectrum covariance matrix thus has a rank of 6. Since there are two components in the sky model (let us call their associated HILC maps \hat{T} and \hat{y}), the largest possible HILC data vector is $\{C_{\ell}^{\hat{T}\hat{T}}, C_{\ell}^{\hat{T}\hat{y}}, C_{\ell}^{\hat{y}\hat{y}}\}$. This would result in a HILC power spectrum covariance matrix of rank 3, which is less than the rank of the frequency-frequency power spectrum covariance matrix. As a result, the two methods do not yield identical parameter covariance matrix elements in this scenario. Analytically, it is not obvious which one has a larger covariance, but numerically, we find that the parameter covariance matrices are actually nearly identical in practice (see Appendix D).

4. Three components, three frequencies

Suppose that the three frequency channels are 90, 150, and 220 GHz. Then there are six possible values for each of ij and km : $ij, km \in \{90 \times 90, 150 \times 150, 220 \times 220, 90 \times 150, 90 \times 220, 150 \times 220\}$. The frequency-frequency power spectrum covariance matrix thus has a rank of 6. Let the HILC maps of the three sky components be denoted as

\hat{T} , \hat{y} , and \hat{J} . Then $\hat{p}\hat{q}$ may take on values of all the auto-spectra and cross spectra, i.e., $\hat{p}\hat{q} \in \{C_{\ell}^{\hat{T}\hat{T}}, C_{\ell}^{\hat{y}\hat{y}}, C_{\ell}^{\hat{J}\hat{J}}, C_{\ell}^{\hat{T}\hat{y}}, C_{\ell}^{\hat{T}\hat{J}}, C_{\ell}^{\hat{y}\hat{J}}\}$. This matrix also has a rank of 6 since it is built from the frequency-frequency covariance matrix of rank-6. Thus, in this case we can use all the possible autospectra and cross spectra in the HILC data vector. Because of computational constraints of running this scenario entirely analytically, we assume a Gaussian frequency-frequency power spectrum covariance matrix and derive analytic results for the final Fisher matrices in both the HILC and multifrequency power spectrum approaches. Testing several numerical setups, we then find that the HILC and multifrequency approaches yield identical parameter covariance matrices, as in the two-component, two-frequency case considered above.

5. General conclusions

Based on these simple examples, we can extrapolate some general conclusions:

- (i) When the number of frequencies is equal to the number of components, we can use all possible HILC map autospectra and cross spectra in the data vector for the HILC approach, and the two approaches yield identical results for the final parameter covariance matrix.
- (ii) When there are more components than frequencies, the rank of the HILC power spectrum covariance matrix is limited by the rank of the frequency-frequency power spectrum covariance matrix, and

we must thus limit the number of spectra in the HILC data vector. After doing so, we find that the two approaches also yield equal parameter covariance matrices.

- (iii) When there are more frequencies than components, the rank of the HILC power spectrum covariance matrix is necessarily less than that of the frequency-frequency power spectrum covariance matrix. Thus, the two methods do not give equivalent results analytically, though numerically we find that they do give nearly identical results, with the HILC posterior often being very slightly larger (see Appendix D).

IX. COMPARING POSTERIOR DISTRIBUTIONS FROM SIMULATIONS

In this section, we compare the simulation-derived posteriors obtained from the multifrequency power spectrum, harmonic ILC, and needlet ILC approaches. The posterior distributions for both LFI and the Gaussian likelihood (with posteriors computed via MLE here) are shown in Fig. 5. In both cases, we see that NILC significantly shrinks the area of the 2D posterior compared to the other approaches. In particular, we see an approximately 60% reduction in the area of the 68% confidence interval 2D posterior (in the elliptical approximation, discussed further in Appendix D) using NILC power spectra as compared with multifrequency power spectra as summary statistics for LFI. Figure 6 compares the posteriors from LFI and the Gaussian likelihood for

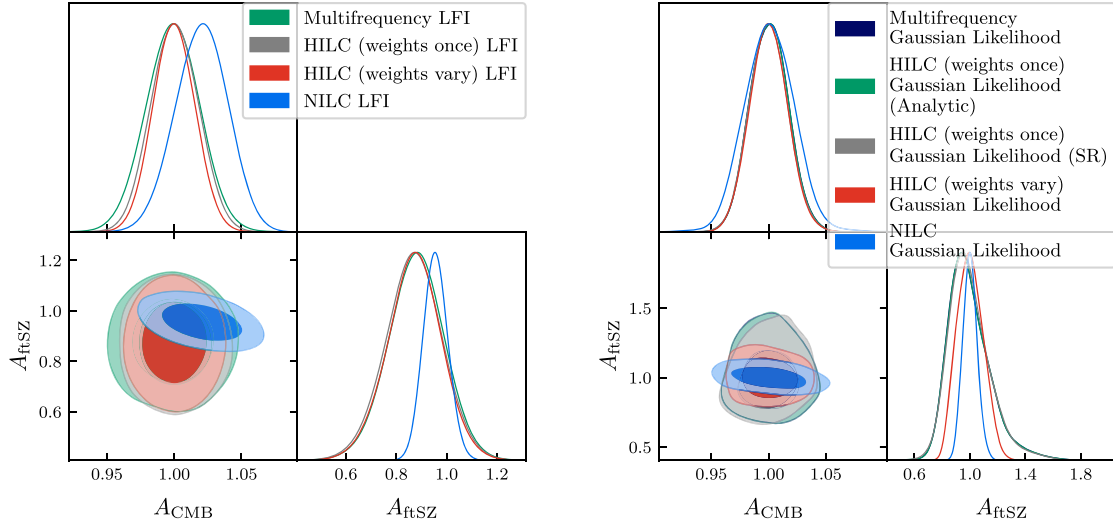


FIG. 5. Parameter posteriors obtained via LFI (left) and Gaussian likelihoods (right). In both figures, we compare the results using various summary statistics: multifrequency power spectra (described in Sec. V), HILC power spectra (with a few variants based on weight determination and parameter dependence determination that are described in Sec. VI), and NILC power spectra (described in Sec. VII). For LFI, we use 40000 simulations for each pipeline. For the Gaussian likelihood, we use 2000 simulations and show the distribution of maximum-likelihood estimates from each simulation. It is evident that the area of the 2D posterior is significantly smaller in the NILC pipeline than the traditional multifrequency pipeline, suggesting that NILC is a useful transformation of the data to use in parameter inference pipelines when there are non-Gaussian sky components.

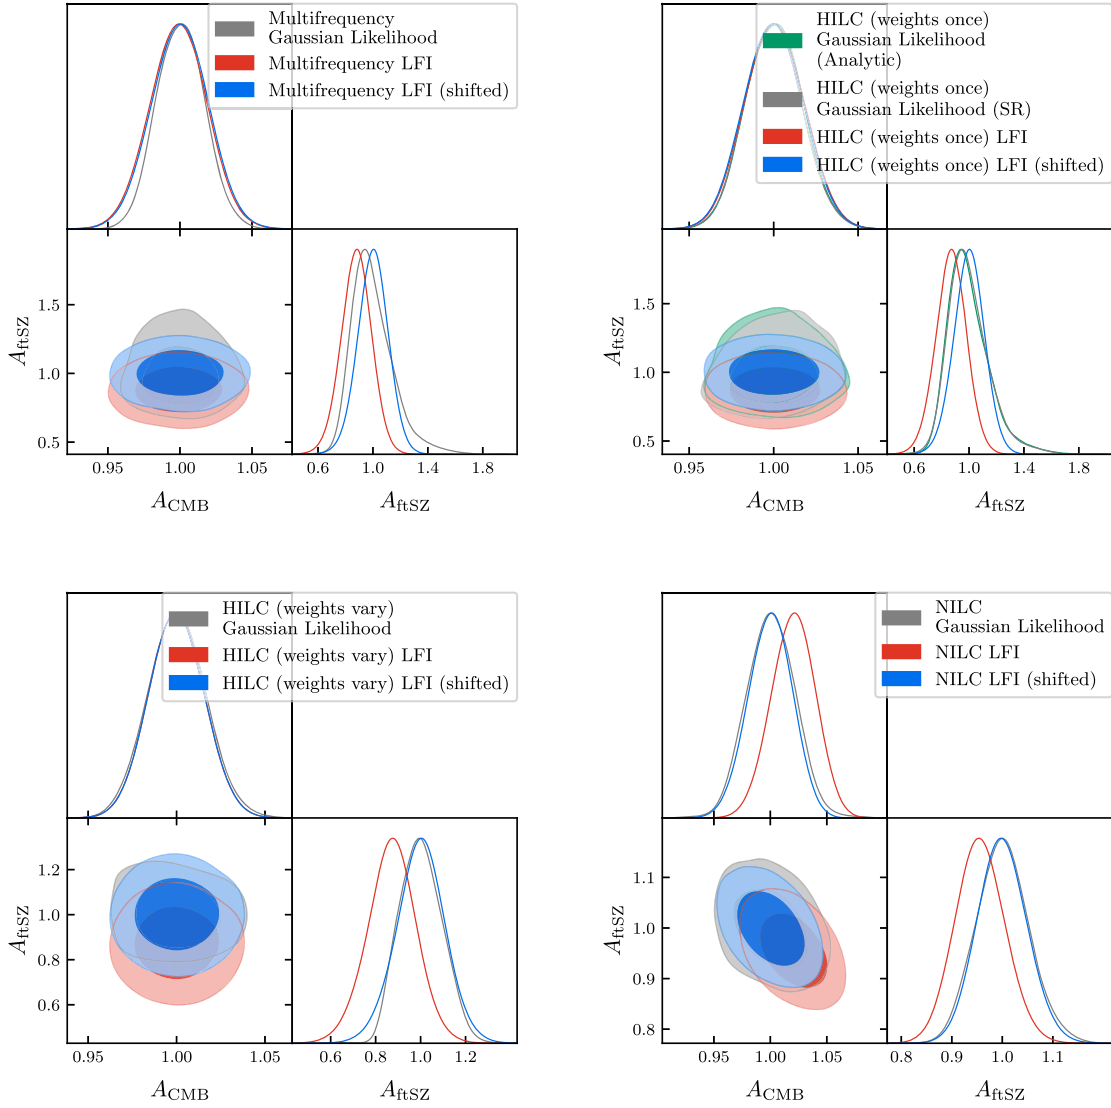


FIG. 6. Comparison of posteriors from LFI to those from a Gaussian likelihood for all the summary statistics considered in this work. We show posteriors for the Gaussian likelihood, LFI, and LFI shifted such that the posterior is centered on 1.0 for both parameters. The Gaussian likelihood is inaccurate in the presence of the highly amplified non-Gaussian tSZ field.

each of the pipelines to assess whether the Gaussian likelihood is a valid assumption. We predict that it is not a good assumption since, if we histogram the value of one summary statistic in one bin from several simulations, the histogram is non-Gaussian. From the changes in the posteriors (computed in the Gaussian likelihood versus with LFI), we see that it is indeed an imperfect assumption. Table II shows the numerical values of 1D marginalized posteriors for all the methods considered in this work.

As a consistency check, we also run all of the pipelines in this work using a setup containing only Gaussian random fields (i.e., we generate Gaussian realizations of the ftSZ field using a fiducial power spectrum). In this situation, there should be no advantage to performing NILC over HILC, and moreover, the Gaussian likelihood should yield

the same results as LFI.⁵ This is explored in Appendix C, where we find that this is indeed the case, and the posteriors from all methods are nearly identical.

As discussed in Sec. VIII B, when the number of frequencies is equal to the number of components in the sky model (as is the case here), the HILC fixed-weights pipeline and multifrequency-power-spectra pipeline must

⁵When data are cosmic-variance limited, LFI is more optimal and less biased than the Gaussian likelihood, even when every field in the problem is Gaussian. This is because the covariance matrix in the Gaussian likelihood is actually a parameter-dependent quantity, but generally treated as parameter independent, while LFI can learn such a dependence [74]. However, in the situation studied here, there is non-negligible instrumental noise, so the data are not cosmic-variance limited, and thus we expect the two approaches to yield nearly identical results.

TABLE II. Parameter constraints for all the methods considered in this work. See Secs. V–VII for information on using multifrequency power spectra, harmonic ILC power spectra, and needlet ILC power spectra as summary statistics, respectively. We present 1D marginalized posterior values obtained with the different summary statistics using both likelihood-free inference and a Gaussian likelihood. For the Gaussian likelihood, we compute posteriors using maximum-likelihood estimation (MLE), Fisher matrix calculations, and a Markov chain Monte Carlo (MCMC) algorithm. Note that we do not quote a central value for the Fisher matrix calculation since it only predicts the error bars (which are assumed to be symmetric in the Fisher approximation).

| Pipeline (on non-Gaussian simulations) | A_{CMB} | A_{rSZ} |
|--|---------------------------------|---------------------------------|
| Multifrequency PS | MLE: 1.001 ± 0.017 | MLE: $0.999_{-0.18}^{+0.10}$ |
| -Gaussian likelihood | Fisher: $_\pm 0.017$ | Fisher: $_\pm 0.16$ |
| -Analytic parameter dependence | MCMC: 1.000 ± 0.016 | MCMC: 0.99 ± 0.16 |
| Multifrequency PS | 0.999 ± 0.020 | 0.88 ± 0.11 |
| -Likelihood-free inference | | |
| Harmonic ILC | MLE: 1.001 ± 0.017 | MLE: $0.999_{-0.18}^{+0.10}$ |
| -Fixed weights | Fisher: $_\pm 0.017$ | Fisher: $_\pm 0.16$ |
| -Gaussian likelihood | MCMC: $1.000_{-0.015}^{+0.017}$ | MCMC: 1.01 ± 0.16 |
| -Analytic parameter dependence | | |
| Harmonic ILC | MLE: 1.000 ± 0.017 | MLE: $1.00_{-0.17}^{+0.10}$ |
| -Fixed weights | Fisher: $_\pm 0.017$ | Fisher: $_\pm 0.16$ |
| -Gaussian likelihood | MCMC: 0.9996 ± 0.017 | MCMC: 0.999 ± 0.16 |
| -Parameter dependence from symbolic regression | | |
| Harmonic ILC | 1.001 ± 0.017 | 0.87 ± 0.11 |
| -Fixed weights | | |
| -Likelihood-free inference | | |
| Harmonic ILC | MLE: 1.000 ± 0.016 | MLE: $1.001_{-0.10}^{+0.089}$ |
| -Varying weights | Fisher: $_\pm 0.016$ | Fisher: $_\pm 0.11$ |
| -Gaussian likelihood | MCMC: 1.002 ± 0.016 | MCMC: $1.035_{-0.11}^{+0.073}$ |
| -Parameter dependence from symbolic regression | | |
| Harmonic ILC | 1.000 ± 0.015 | 0.87 ± 0.11 |
| -Varying weights | | |
| -Likelihood-free inference | | |
| Needlet ILC | MLE: 1.001 ± 0.022 | MLE: 1.001 ± 0.055 |
| -Gaussian likelihood | Fisher: $_\pm 0.022$ | Fisher: $_\pm 0.052$ |
| -Parameter dependence from symbolic regression | MCMC: 1.001 ± 0.022 | MCMC: $1.007_{-0.054}^{+0.049}$ |
| Needlet ILC | 1.020 ± 0.019 | 0.956 ± 0.049 |
| -Likelihood-free inference | | |

yield identical results for the Gaussian-likelihood analysis. We have shown that when the number of components is greater than the number of frequency channels, the fixed-weights HILC power spectrum covariance matrix becomes singular if we make HILC maps for every component in the sky model and use all of their autospectra and cross-power spectra in the data vector. A natural question is whether the same would be true for NILC power spectra, which have nontrivial scale-dependent coupling. Using our simulations, we show that, to within our numerical precision, the determinant of the NILC power spectrum covariance matrix

is zero in such a situation. Specifically, we demonstrate this on a setup very similar to the main one used in the simulations, but with an added Gaussian realization of the CIB field, assuming a modified blackbody SED [75] with an effective dust temperature of 20.0 K and spectral index of 1.45. With this setup, we also verify numerically that the HILC power spectrum covariance matrices are singular for both variants of the HILC pipeline (fixed weights and varying weights). The implication is that, for both the HILC and NILC inference approaches, if the number of components is greater than the number of frequency channels, one

must carefully construct the fundamental data vector such that the elements are linearly independent (i.e., one should not use the autospectra and cross spectra of every component ILC map in the data vector).

We also examine how the results change when varying the frequency channels used or when increasing the number of frequency channels to three (while still keeping just two components). These results are shown in Appendix D. In all cases, using NILC power spectra as summary statistics either shrinks or yields the same area of the 2D posterior as compared with using multifrequency power spectra as summary statistics. This method thus has significant promise in yielding tighter parameter constraints from ongoing and upcoming CMB experiments.

X. DISCUSSION

In this paper we have considered several methods for inferring parameters describing the power spectra of physical sky components from multifrequency mm-wave data. We have shown that using autospectra and cross spectra of HILC maps in a Gaussian likelihood yields the same final parameter error bars as the traditional approach of fitting templates to multifrequency autospectra and cross-power spectra, with caveats of which HILC spectra to include in the data vector based on the number of frequency channels and number of sky components (discussed in Sec. VIII). Specifically, when the number of sky components is greater than or equal to the number of frequency channels, we have analytically shown that performing inference with harmonic ILC power spectra (with fixed weights) yields identical parameter covariance matrices as performing inference with multifrequency power spectra on Gaussian random fields. However, the power spectra of HILC maps only contain Gaussian information, and many mm-wave foreground fields vary spatially on the sky and are non-Gaussian. Thus, NILC is a better candidate to capture such additional information. In previous work, we derived an analytic expression for the autospectra and cross spectra of NILC maps in Ref. [24], finding a nontrivial parameter dependence. In this work, we have estimated this parameter dependence from simulations using symbolic regression, and have then performed a maximum-likelihood analysis to show that NILC reduces the area of the 68% confidence interval 2D posterior on CMB and ftSZ amplitude parameters by 60%, with results summarized in Table II and 2D and 1D marginalized posterior distributions shown in Fig. 5.

Because the weight maps are functions of the fields, the results of Ref. [24] demonstrate that the parameter dependence of NILC power spectra is highly nontrivial. With NILC, the exact dependence of the weight maps on the fields is complex, so we must estimate this dependence

using simulations and symbolic regression to avoid assuming a specific functional form. Figure 10 in Appendix B shows the complexity of this dependence, which depends on the exact amplitude of each component. In some cases, we note that when the amplitude of some component increases, its contribution to the NILC power spectrum of another component actually decreases since the weight maps then prioritize suppressing this contaminant over the others. With such complex parameter dependence, and given that in many situations the Gaussian likelihood is an imperfect assumption, LFI is a much simpler and more accurate approach for determining parameter posteriors in the presence of non-Gaussian contaminants. Similar to the case when using the Gaussian likelihood, we find that performing LFI with NILC power spectra as summary statistics reduces the area of the 68% confidence interval 2D posterior on CMB and ftSZ amplitude parameters by 60% as compared with performing LFI with multifrequency power spectra as summary statistics.

In Appendix D we examine various frequency combinations and how they impact the results of the different summary statistics. In all cases, NILC either maintains or shrinks the area of the 2D posterior, showing that NILC component map power spectra are powerful summary statistics. In our setup, we have used 40,000 simulations in the LFI, using a single round of NPE with a uniform prior. To be more efficient in the number of simulations, one could use sequential neural posterior estimation (SNPE), running the NPE procedure in multiple rounds [50–53]. In the first round, parameters are sampled from the prior and used to generate simulations, learning an initial posterior. In subsequent rounds, parameters are sampled instead from the current learned posterior rather than the prior. Another modification to be more efficient in the number of simulations is to use a very wide Gaussian prior rather than a uniform prior.

Here we have demonstrated our result on a somewhat contrived scenario in which the tSZ field is amplified by a factor of 150, but we would expect similar trends in the error bars with the usual tSZ field at high ℓ , in particular, beyond $\ell = 2000$ where the tSZ signal amplitude surpasses that of the CMB at 90 GHz. Moreover, we would expect such large error bar reductions in other cases where non-Gaussian foregrounds dominate over the signal of interest, as is the case when searching for primordial B modes, for which the non-Gaussian dust dominates. With the latest constraint on the tensor-to-scalar ratio of $r_{0.05} < 0.036$ [5], the dust amplitude is approximately a factor of 10–100 times larger than the CMB B -mode amplitude at 150 GHz for $2 \leq \ell \leq 500$ (see Fig. 12 of Ref. [76]), depending on the region of sky analyzed. This ratio is even larger than the ratio in our toy model of the ftSZ to CMB (which is within a factor of a few, cf. Fig. 1). The current limitation in

demonstrating our NILC approach on the example of primordial B modes and polarized dust is the need for many realistic anisotropic, non-Gaussian simulations of the dust field. However, progress has been made toward such simulations in recent work [77–79], and is also the focus of follow-up work in preparation.

This work demonstrates the ability of NILC to capture non-Gaussian information, as well as the constraining power that such information can have on parameters. It thus motivates further methodological development for using NILC-based inference or other techniques that efficiently capture information beyond two-point statistics in analyses of the mm-wave sky.

ACKNOWLEDGMENTS

We thank Marcelo Alvarez for help with generating tSZ simulations with HALOSKY; and Adrian Bayer, Sam Goldstein, Mathew Madhavacheril, Fiona McCarthy, Chirag Modi, Shivam Pandey, Oliver Philcox, and Kendrick Smith for useful discussions. We also thank Adri Duivenvoorden for comments on the manuscript. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE 2036197 (K. M. S.). J. C. H. acknowledges support from NSF Grant No. AST-2108536, NASA Grant No. 80NSSC22K0721, NASA Grant No. 80NSSC23K0463, DOE Grant No. DE-SC00233966, the Sloan Foundation, and the Simons Foundation. Several software tools were used in the development and presentation of results shown in this paper, including GetDist [80], HEALPix/HEALPY [61,62], NumPy [81], SciPy [82], Matplotlib [83], Astropy [84–86], EMCEE [68], PySR [71], SBI [69], PYILC [17,36], WANDB [70], PyTorch [87], and Julia [88]. The authors acknowledge the Texas Advanced Computing Center (TACC) at The University of Texas at Austin for providing HPC resources that have contributed to the research results reported within this paper. The authors also acknowledge the use of resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory. Finally, the authors acknowledge the use of computing resources from Columbia University’s Shared Research Computing Facility project, which is supported by NIH Research Facility Improvement Grant No. 1G20RR030893-01, and associated funds from the New York State Empire State Development, Division of Science Technology and Innovation (NYSTAR) Contract No. C090171, both awarded April 15, 2010.

APPENDIX A: ILC MAPS AND WEIGHT MAPS

In this appendix we show images of component maps and ILC weight maps. For the harmonic ILC pipeline, we compute the spectra that comprise the data vector analytically, as described in Sec. VI. However, one can also explicitly construct a HILC map and then compute the power spectrum of that map. In Fig. 7, we compare the analytic spectra to the map-based spectra, confirming that the results are effectively identical.

In Fig. 8, we compare the reconstructed CMB and amplified Compton- y maps from the NILC and HILC pipelines. The reconstructed amplified Compton- y maps look very similar. However, in the reconstructed CMB maps, one can see some bright spots in the HILC map that are not present in the input CMB map or NILC CMB map. These coincide with bright spots in the ftSZ map. Thus, this is an example of how HILC does not optimally clean out non-Gaussian contaminants, but NILC does.

Figure 9 shows the NILC weight maps used for the construction of both the CMB-preserved NILC map and ftSZ-preserved NILC map. These weight maps are built from the same map realizations shown in Fig. 8. We see that the weight maps have the general appearance of the contaminant field since the purpose of the weight maps is to suppress contamination. For example, the weight maps for a CMB-preserved NILC map roughly mimic the structure of the ftSZ map in Fig. 8.

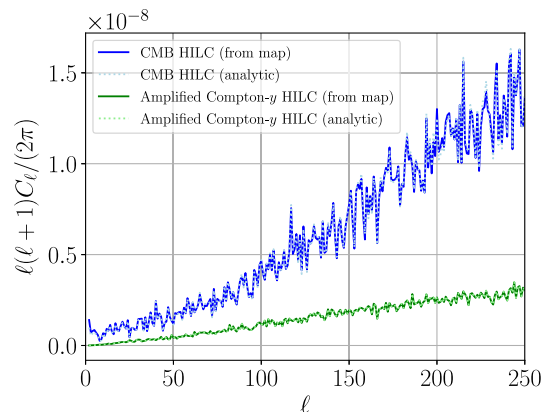


FIG. 7. Comparison of harmonic ILC spectra computed analytically as described in Sec. VI (dotted curves) and spectra computed from a map-based HILC operation using PYILC (solid curves) for a single simulation. The curves are shown in units of K^2 for the CMB spectra and dimensionless Compton- y units for the Compton- y spectra. Harmonic ILC weights are computed with $\Delta\ell = 20$ (see Sec. VI). There is no ℓ -space binning shown in the plot.

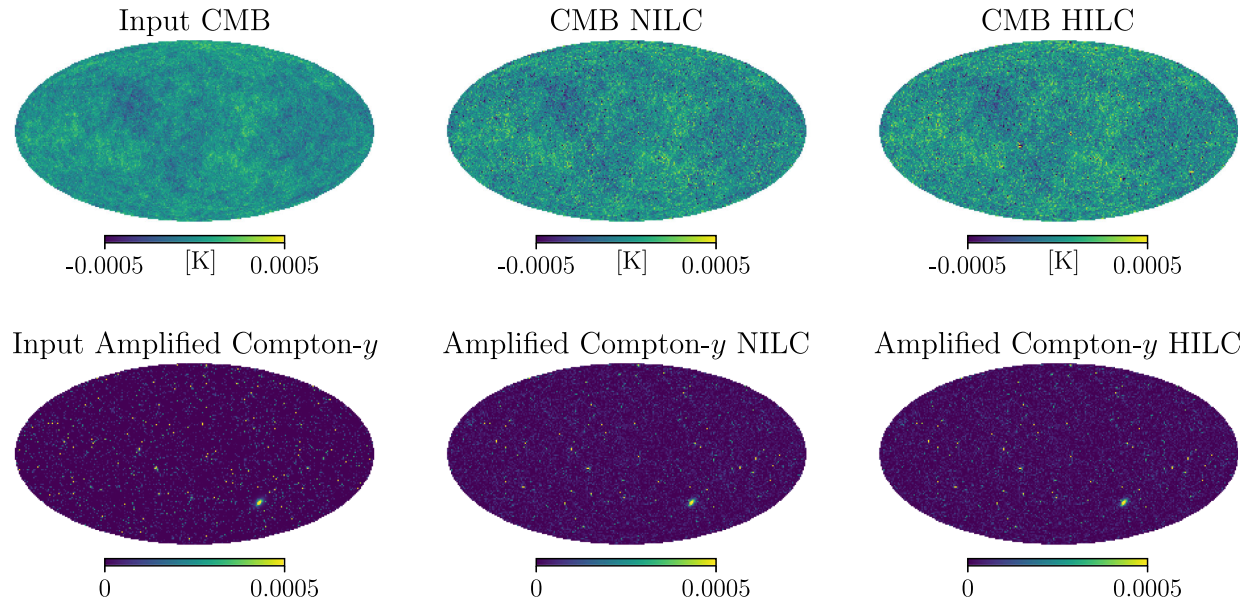


FIG. 8. Comparison of input CMB and amplified Compton- y maps to the maps reconstructed via NILC and HILC. Residual bright spots resulting from ftSZ contamination are present in the CMB HILC map but not in the CMB NILC map, demonstrating the ability of NILC to better clean non-Gaussian contaminants.

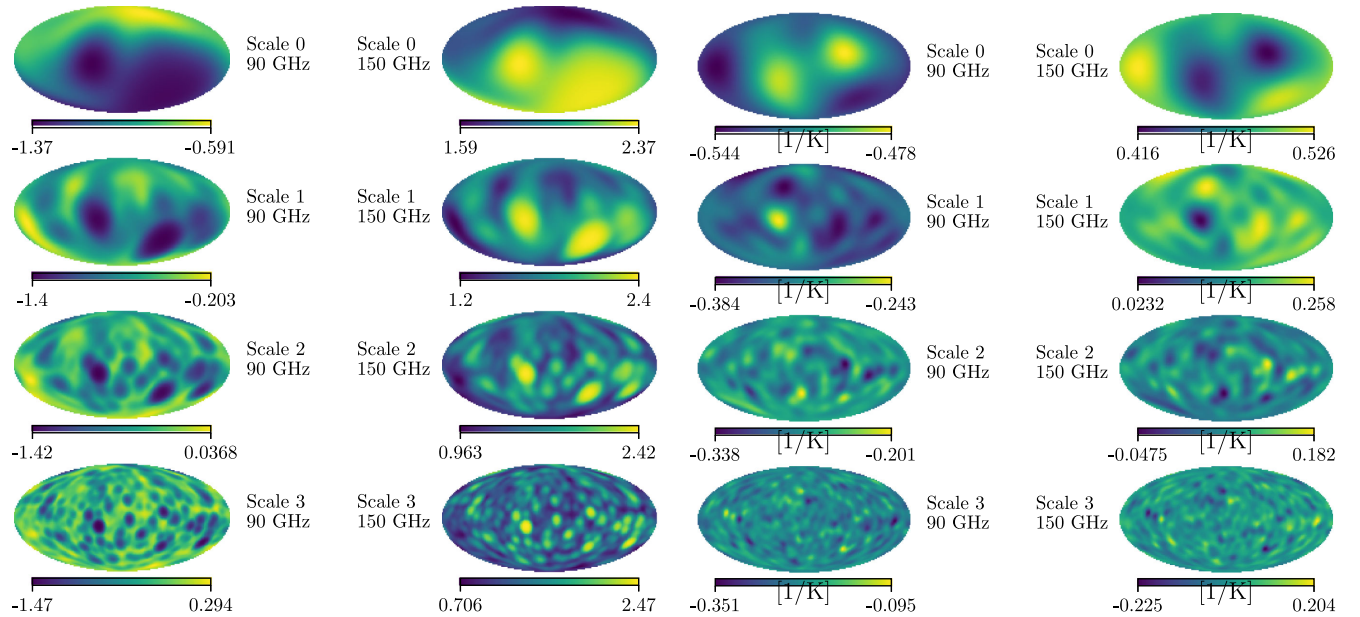


FIG. 9. NILC weight maps for a preserved CMB component (left two columns) and preserved tSZ component (right two columns), generated with PYILC, for a sky model consisting of the CMB, amplified tSZ signal (ftSZ), and noise. The maps are shown for 90 and 150 GHz at each needlet filter scale, with scale 0 corresponding to the lowest- ℓ scale and scale 3 corresponding to the highest- ℓ scale. The CMB and ftSZ maps underlying the frequency maps from which these NILC weight maps are produced are shown in Fig. 8. The CMB weight maps are unitless, and the Compton- y weight maps have units of $1/\text{K}$ such that the output NILC map is in dimensionless Compton- y units.

APPENDIX B: PARAMETER DEPENDENCE FROM SYMBOLIC REGRESSION

In this appendix we discuss results for determining the parameter dependence of HILC and NILC power spectra via symbolic regression. The details of this procedure are presented in Sec. VI A 2. In Fig. 10 we show the values of $f_b^{\hat{p}\hat{q}}(A_{\text{CMB}}, A_{\text{rSZ}})$ from Eq. (27) for the different \hat{p}, \hat{q} and different methods considered in this work. Since f_b is a separate function for each bin, we choose one bin for which to plot the results (here it is $b = 5$, which

corresponds to a mean multipole of $\ell = 138$). We show the values of f_b for the region of parameter space around the fiducial $A_{\text{CMB}} = 1 = A_{\text{rSZ}}$.

We also show the analytic expressions output by the regressor for bin 5 in Table III. From these expressions, we see that the regressor may be slightly overfitting (for example, this is apparent from the term involving A_{rSZ} in the expression for NILC $C_b^{\hat{T}\hat{T}}$), but this is not important for our purposes here. We simply require expressions that give smooth and accurate parameter dependence around the fiducial $A_{\text{CMB}} = 1 = A_{\text{rSZ}}$.




FIG. 10. Visual representation of the parameter dependence of ILC component map power spectra obtained via symbolic regression for NILC power spectra (top row), HILC power spectra with weights determined separately for each simulation (middle row), and HILC power spectra with weights determined from a fiducial template and then applied to every simulation (bottom row). Parameter dependence is shown for the different elements of the data vector: $C_\ell^{\hat{T}\hat{T}}$ (left), $C_\ell^{\hat{T}\hat{y}}$ (middle), and $C_\ell^{\hat{y}\hat{y}}$ (right). In particular, we plot $f_b^{\hat{p}\hat{q}}(A_{\text{CMB}}, A_{\text{rSZ}})$ from Eq. (27), where $b = 5$ here, corresponding to a mean multipole $\ell = 138$.

TABLE III. Parameter dependence of ILC power spectra obtained via symbolic regression for NILC power spectra, HILC power spectra with weights determined separately for each simulation, and HILC power spectra with weights determined from a fiducial template and then applied to every simulation. Parameter dependence is shown for the different elements of the data vector: $C_{\ell}^{\hat{T}\hat{T}}$, $C_{\ell}^{\hat{T}\hat{y}}$, and $C_{\ell}^{\hat{y}\hat{y}}$. In particular, we show the resulting expressions for $f_b^{\hat{p}\hat{q}}(A_{\text{CMB}}, A_{\text{ftSZ}})$ from Eq. (27), where $b = 5$ here, corresponding to a mean multipole $\ell = 138$.

| Spectrum | Expression |
|---|---|
| NILC $C_{b=5}^{\hat{T}\hat{T}}$ | $0.57870084A_{\text{CMB}} - 1.53680702801541 \times 10^{-5}(A_{\text{ftSZ}})^{27} + 0.42053628$ |
| NILC $C_{b=5}^{\hat{T}\hat{y}}$ | $\exp\left(\frac{A_{\text{CMB}} - 1/A_{\text{ftSZ}}}{A_{\text{CMB}}/A_{\text{ftSZ}} + 2.737616}\right)$ |
| NILC $C_{b=5}^{\hat{y}\hat{y}}$ | $0.03574537A_{\text{CMB}} - 0.03574537(A_{\text{ftSZ}})^2 + A_{\text{ftSZ}}$ |
| HILC (weights vary) $C_{b=5}^{\hat{T}\hat{T}}$ | $\exp\{[(A_{\text{CMB}})^4 - A_{\text{ftSZ}}] \times \exp(-2A_{\text{CMB}})\}$ |
| HILC (weights vary) $C_{b=5}^{\hat{T}\hat{y}}$ | $\exp\{0.22931802A_{\text{CMB}} - 0.22931802 \times \frac{\exp[(A_{\text{CMB}} - A_{\text{ftSZ}})^3]}{A_{\text{ftSZ}}}\}$ |
| HILC (weights vary) $C_{b=5}^{\hat{y}\hat{y}}$ | $0.036737457A_{\text{CMB}} - 0.036737457(A_{\text{ftSZ}})^2 + A_{\text{ftSZ}}$ |
| HILC (weights fixed) $C_{b=5}^{\hat{T}\hat{T}}$ | $0.5588486A_{\text{CMB}} + 0.4411514A_{\text{ftSZ}}$ |
| HILC (weights fixed) $C_{b=5}^{\hat{T}\hat{y}}$ | $0.24519494A_{\text{CMB}} + 0.75480506A_{\text{ftSZ}}$ |
| HILC (weights fixed) $C_{b=5}^{\hat{y}\hat{y}}$ | $0.076038094388331A_{\text{CMB}} + 0.923961905611669A_{\text{ftSZ}}$ |

APPENDIX C: RESULTS ON GAUSSIAN RANDOM FIELDS

In this appendix we present the same results as in Sec. IX, but evaluated on simulations containing only Gaussian random fields. The CMB and noise remain the same as described in Sec. IV, but the ftSZ field is now a Gaussian random field as well. Specifically, we take the

power spectrum of one tSZ map generated via HALOSKY, where the map is amplified by a factor of 150. Then for each simulation, we generate a Gaussian realization of that power spectrum. With all Gaussian random fields, we expect that the posterior distributions should be nearly identical for all of the methods considered in this work. In particular, the Gaussian likelihood should be equivalent to

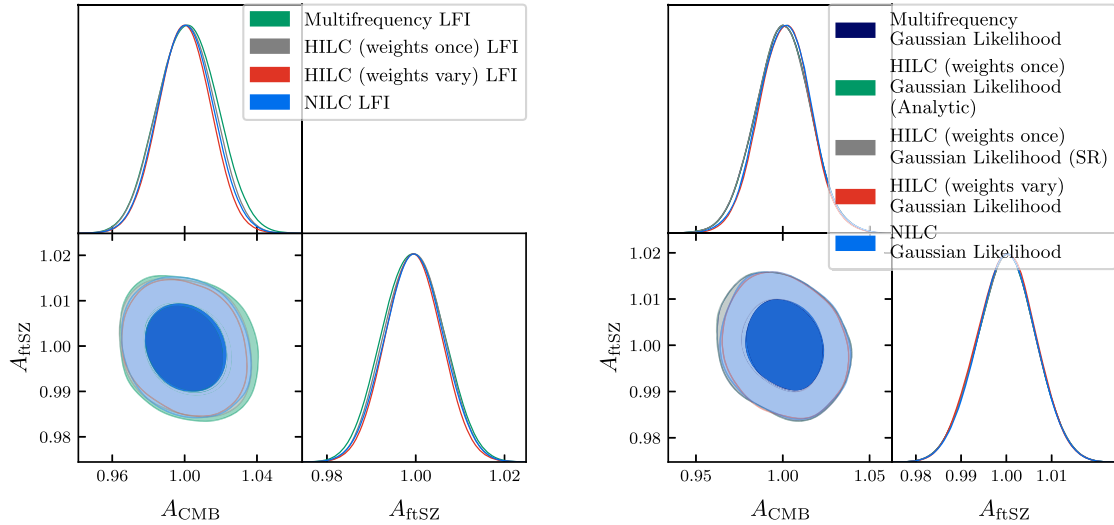


FIG. 11. Parameter posteriors obtained via LFI (left) and Gaussian likelihoods (right) using simulations containing only Gaussian random fields (here the ftSZ field is a Gaussian realization of an amplified template tSZ power spectrum). In both figures, we compare the results using various summary statistics: multifrequency power spectra (described in Sec. V), harmonic ILC (HILC) power spectra (with a few variants based on weight determination and parameter dependence determination that are described in Sec. VI), and NILC power spectra (described in Sec. VII). For LFI, we use 30000 simulations for each pipeline. For the Gaussian likelihood, we use 2000 simulations and show the distribution of maximum-likelihood estimates from each simulation.

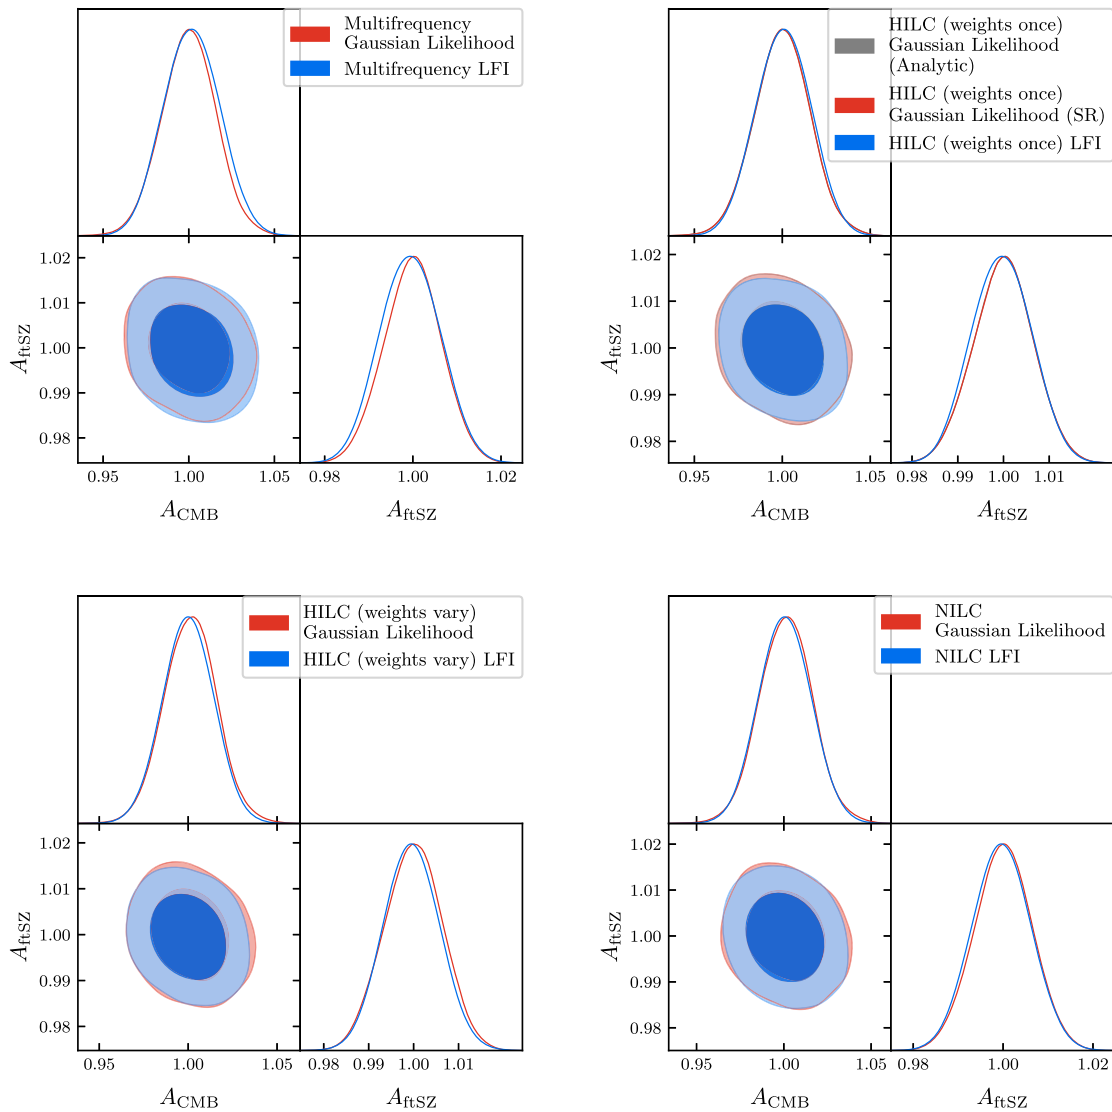


FIG. 12. Comparison of posteriors from LFI to those from a Gaussian likelihood for all the summary statistics considered in this work, computed on simulations consisting only of Gaussian random fields. The Gaussian likelihood is a reasonable assumption when every field in the simulation is a Gaussian random field.

likelihood-free inference since the Gaussian likelihood assumption must hold when every field in the problem is Gaussian random.⁶ Moreover, without the presence of large non-Gaussian fields, we would expect that there is no additional information for NILC to extract that is not already encompassed by the two-point functions contained in the multifrequency power spectra. Thus, we also expect that using multifrequency power spectra, HILC power spectra, and NILC power spectra as summary statistics should all yield nearly identical results. We use this fact to

⁶Strictly speaking, this is only true for sufficiently high ℓ —at low ℓ there can still be deviations from Gaussianity in the likelihood due to the small number of harmonic modes, as described in, e.g., Ref. [89].

validate all of the methods presented in this paper and to assess the significance of 2D posterior area reductions in the full non-Gaussian ftSZ field case.

The posterior distributions for both LFI and the Gaussian likelihood (with posteriors computed via MLE here) are shown in Fig. 11. From these plots, we see that multifrequency, HILC, and NILC power spectra yield nearly identical posteriors. Figure 12 compares the posteriors from LFI and the Gaussian likelihood for each of the pipelines, validating the LFI procedure by showing that it recovers the same results as the Gaussian likelihood in this situation where we know that the Gaussian likelihood should hold. Table IV shows the 1D marginalized posteriors for each of the methods discussed in this work, evaluated on this setup of only Gaussian random fields.

TABLE IV. Parameter constraints for all the methods considered in this work, as applied to simulations containing only Gaussian random fields (here the ftSZ field is a Gaussian realization of an amplified tSZ power spectrum template). See Secs. V–VII for information on using multifrequency power spectra, harmonic ILC power spectra, and needlet ILC power spectra as summary statistics, respectively. We present 1D marginalized posterior values obtained with the different summary statistics using both likelihood-free inference (with 30000 simulations) and a Gaussian likelihood (with 2000 simulations). For the Gaussian likelihood, we compute posteriors using maximum-likelihood estimation (MLE), Fisher matrix calculation, and a Markov chain Monte Carlo (MCMC) algorithm. Note that we do not quote a central value for the Fisher matrix calculation since it only predicts the error bars (which are assumed to be symmetric in the Fisher calculation).

| Pipeline (using only Gaussian random fields) | A_{CMB} | A_{ftSZ} |
|--|--------------------------|----------------------------|
| Multifrequency PS | MLE: 1.000 ± 0.016 | MLE: 0.99998 ± 0.0064 |
| -Gaussian likelihood | Fisher: $_\pm 0.016$ | Fisher: $_\pm 0.0065$ |
| -Analytic parameter dependence | MCMC: 0.999 ± 0.015 | MCMC: 0.99997 ± 0.0064 |
| Multifrequency PS | 1.002 ± 0.016 | 0.9994 ± 0.0068 |
| -Likelihood-free inference | | |
| Harmonic ILC | MLE: 1.000 ± 0.016 | MLE: 0.99998 ± 0.0064 |
| -Fixed weights | Fisher: $_\pm 0.016$ | Fisher: $_\pm 0.0065$ |
| -Gaussian likelihood | MCMC: 0.999 ± 0.016 | MCMC: 1.0001 ± 0.0064 |
| -Analytic parameter dependence | | |
| Harmonic ILC | MLE: 1.000 ± 0.016 | MLE: 1.0000 ± 0.0064 |
| -Fixed weights | Fisher: $_\pm 0.016$ | Fisher: $_\pm 0.0065$ |
| -Gaussian likelihood | MCMC: 0.9997 ± 0.016 | MCMC: 1.0006 ± 0.0062 |
| -Parameter dependence from symbolic regression | | |
| Harmonic ILC | 1.001 ± 0.016 | 0.9996 ± 0.0065 |
| -Fixed weights | | |
| -Likelihood-free inference | | |
| Harmonic ILC | MLE: 1.001 ± 0.015 | MLE: 0.99995 ± 0.0064 |
| -Varying weights | Fisher: $_\pm 0.015$ | Fisher: $_\pm 0.0064$ |
| -Gaussian likelihood | MCMC: 1.002 ± 0.015 | MCMC: 0.9995 ± 0.0064 |
| -Parameter dependence from symbolic regression | | |
| Harmonic ILC | 1.000 ± 0.014 | 0.9995 ± 0.0062 |
| -Varying weights | | |
| -Likelihood-free inference | | |
| Needlet ILC | MLE: 1.001 ± 0.015 | MLE: 1.0000 ± 0.0064 |
| -Gaussian likelihood | Fisher: $_\pm 0.015$ | Fisher: $_\pm 0.0064$ |
| -Parameter dependence from symbolic regression | MCMC: 1.000 ± 0.016 | MCMC: 1.0002 ± 0.0064 |
| Needlet ILC | 1.001 ± 0.015 | 0.9997 ± 0.0065 |
| -Likelihood-free inference | | |

APPENDIX D: EFFECTS OF VARYING FREQUENCY CHANNELS

In Sec. IX we consider a toy model in which we use two frequency channels: 90 and 150 GHz. In this appendix, we consider various other combinations of frequency channels to forecast the impact of NILC on parameter constraints. Our physical sky model is identical to that in Sec. IX, comprising the CMB and ftSZ fields, but we evaluate these at different sets of frequencies here. In particular, we consider the following sets of channels (all in GHz): 90, 120; 90, 150; 90, 280; 90, 120, 150; 90, 150, 280; 90, 150,

220; 280, 353; and 280, 353, 400. The noise models and beams are the same in each frequency channel. We compute the results for both the Gaussian likelihood and likelihood-free inference using each summary statistic (multifrequency power spectra, HILC with fixed weights, HILC with varying weights, and NILC), for each set of frequencies. We then estimate the area of the resulting 2D posteriors by extracting the covariance matrix and assuming an elliptical posterior (a crude approximation but sufficient for the rough estimation performed here). As described in Ref. [90], the ellipse parameters are then given by

TABLE V. Area of 2D posteriors multiplied by a factor of 1000 for clarity (assuming elliptical posteriors) for various frequency combinations, using likelihood-free inference. We note that the elliptical approximation is crude in some cases.

| Frequencies (GHz) | Multifrequency PS | HILC (weights once) | HILC (weights vary) | NILC |
|-------------------|-------------------|---------------------|---------------------|------|
| 90, 120 | 36.0 | 35.2 | 30.6 | 18.8 |
| 90, 150 | 16.0 | 14.2 | 12.2 | 6.3 |
| 90, 280 | 5.5 | 5.3 | 5.1 | 5.0 |
| 90, 120, 150 | 16.0 | 13.4 | 11.9 | 7.2 |
| 90, 150, 280 | 5.3 | 5.4 | 5.0 | 5.3 |
| 90, 150, 220 | 6.0 | 6.1 | 5.7 | 4.6 |
| 280, 353 | 11.4 | 10.3 | 8.5 | 4.6 |
| 280, 353, 400 | 10.8 | 8.3 | 8.2 | 5.3 |

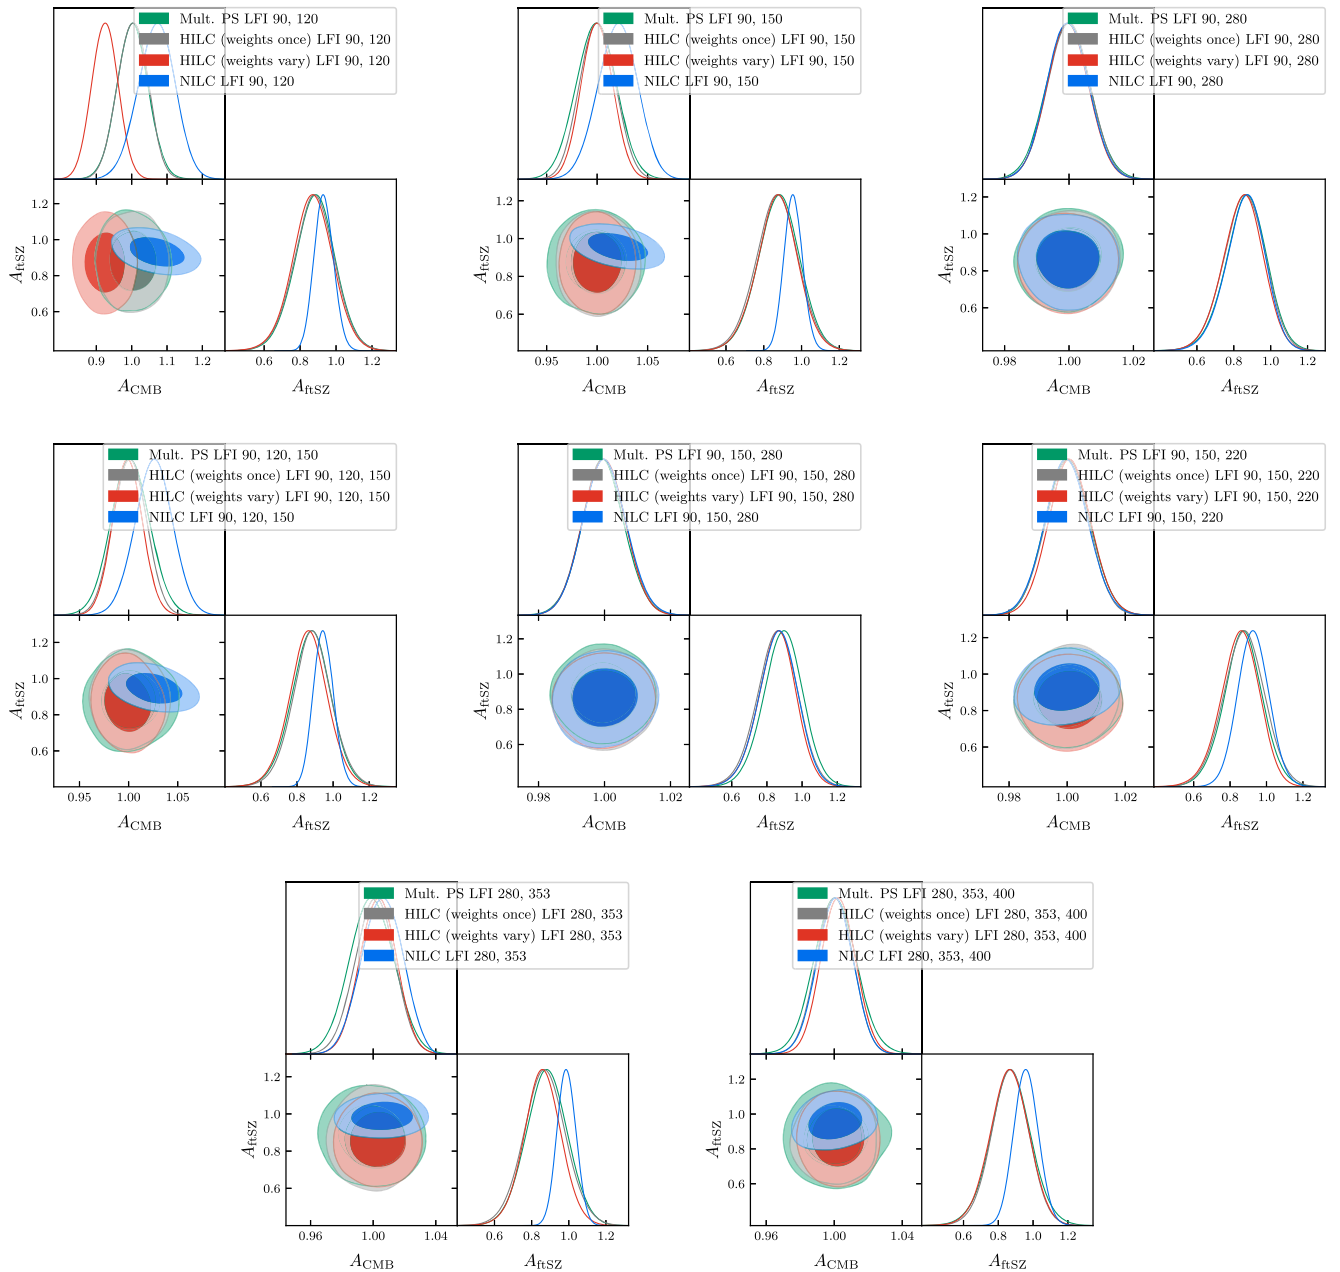


FIG. 13. Posteriors using various summary statistics and frequency combinations with LFI.

$$a^2 = \frac{\sigma_x^2 + \sigma_y^2}{2} + \sqrt{\frac{(\sigma_x^2 - \sigma_y^2)^2}{4} + \sigma_{xy}^2}, \quad (\text{D1})$$

$$b^2 = \frac{\sigma_x^2 + \sigma_y^2}{2} - \sqrt{\frac{(\sigma_x^2 - \sigma_y^2)^2}{4} + \sigma_{xy}^2}, \quad (\text{D2})$$

and the area of ellipse is then

$$A = \pi ab \Delta\chi^2, \quad (\text{D3})$$

where $\Delta\chi^2 = 2.3$ for the 68% confidence interval.

Table V details the 68% confidence ellipse areas for each of the frequency sets for each summary statistic, using likelihood-free inference, and Fig. 13 shows the posteriors, comparing the various summary statistics for each individual set of frequencies. Notably, the 2D posterior obtained using NILC power spectra as a summary statistic (blue contours in Fig. 13) is the same or smaller than the 2D posteriors obtained using multifrequency power spectra as a summary statistic (green contours in Fig. 13) for every set of frequencies. Table VI and Fig. 14 show the analogous information obtained using a Gaussian likelihood instead of LFI.

Using harmonic ILC with fixed weights, ‘‘HILC (weights once),’’ often yields smaller 2D posterior areas than using multifrequency power spectra in the LFI case. From an information perspective, both methods contain the exact same information, as the former just rescales the latter by

some fixed constants that are applied to every simulation in the same way. This indicates that the LFI may not yet have converged for the multifrequency power spectrum case. Thus, this simple rescaling into HILC power spectra may aid in the algorithm’s ability to learn parameters by transforming the data into a space that is more immediately connected to the parameters of interest.

Also of note is that, in most cases, adding additional frequency channels increases the constraining power of a given summary statistic, where the constraining power is estimated by the area of the Fisher ellipse. As can be seen in Table V, there are a few exceptions to this trend. From an information content standpoint, adding frequency channels can only add information and thus increase constraining power. Thus, we attribute minor deviations from the expected trend to lack of full LFI convergence (or need for more simulations in the Gaussian likelihood case) and the crude approximation that the 2D posteriors are elliptical.

In our setup, we have selected the amplification factor of the Compton- γ field specifically such that the resulting ftSZ power spectrum would have comparable magnitude to that of the CMB. Thus, depending on the frequency used, the tSZ SED could make the ftSZ power spectrum either larger or smaller than that of the CMB. This would play a role in which parameter is most well constrained. We observe this effect in the NILC posteriors in Fig. 15, where the orientation of the posterior interestingly flips, for example, when switching from using 90 and 150 GHz channels to using 90, 150, and 280 GHz channels.

TABLE VI. Area of 2D posteriors multiplied by a factor of 1000 for clarity (assuming elliptical posteriors) for various frequency combinations, using Gaussian likelihoods. For the HILC (weights once) case, we use the analytic parameter dependence. We note that the elliptical approximation is crude in some cases. In particular, for the case of 280, 353, and 400 GHz, it appears that the HILC (weights once) posterior is significantly smaller than the multifrequency PS posterior. However, the 1D marginalized posteriors are $A_{\text{CMB}} = 1.000 \pm 0.010$ and $A_{\text{ftSZ}} = 1.00^{+0.10}_{-0.18}$ in the HILC (weights once) pipeline, and $A_{\text{CMB}} = 1.000 \pm 0.010$ and $A_{\text{ftSZ}} = 0.999^{+0.12}_{-0.17}$ in the multifrequency PS case, and thus, the two methods actually produce nearly identical numerical posteriors.

| Frequencies (GHz) | Multifrequency PS | HILC (weights once) | HILC (weights vary) | NILC |
|-------------------|-------------------|---------------------|---------------------|------|
| 90, 120 | 46.8 | 46.5 | 25.4 | 19.1 |
| 90, 150 | 19.5 | 19.5 | 11.2 | 8.3 |
| 90, 280 | 7.9 | 7.9 | 7.8 | 7.9 |
| 90, 120, 150 | 19.1 | 19.4 | 12.4 | 10.2 |
| 90, 150, 280 | 7.5 | 7.7 | 7.8 | 7.7 |
| 90, 150, 220 | 8.5 | 8.7 | 9.0 | 6.6 |
| 280, 353 | 14.4 | 14.4 | 12.3 | 5.8 |
| 280, 353, 400 | 14.2 | 11.8 | 11.7 | 7.6 |

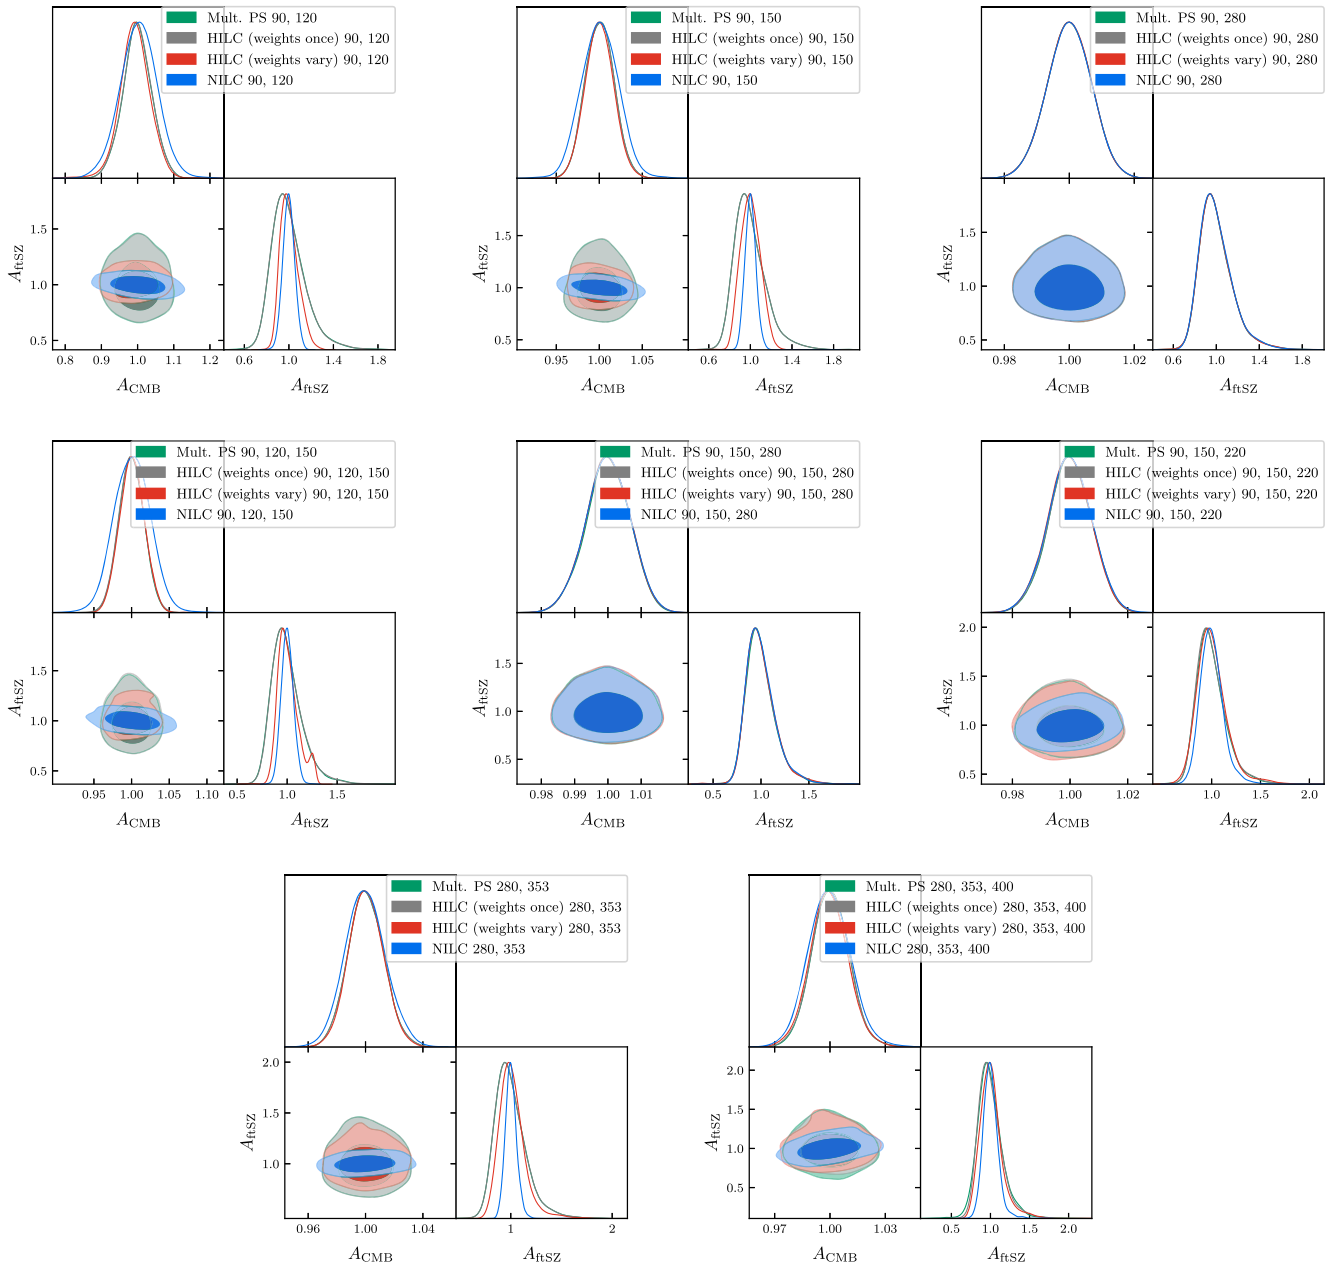


FIG. 14. Posteriors using various summary statistics and frequency combinations with Gaussian likelihoods. For the HILC (weights once) case, we use the analytic parameter dependence.

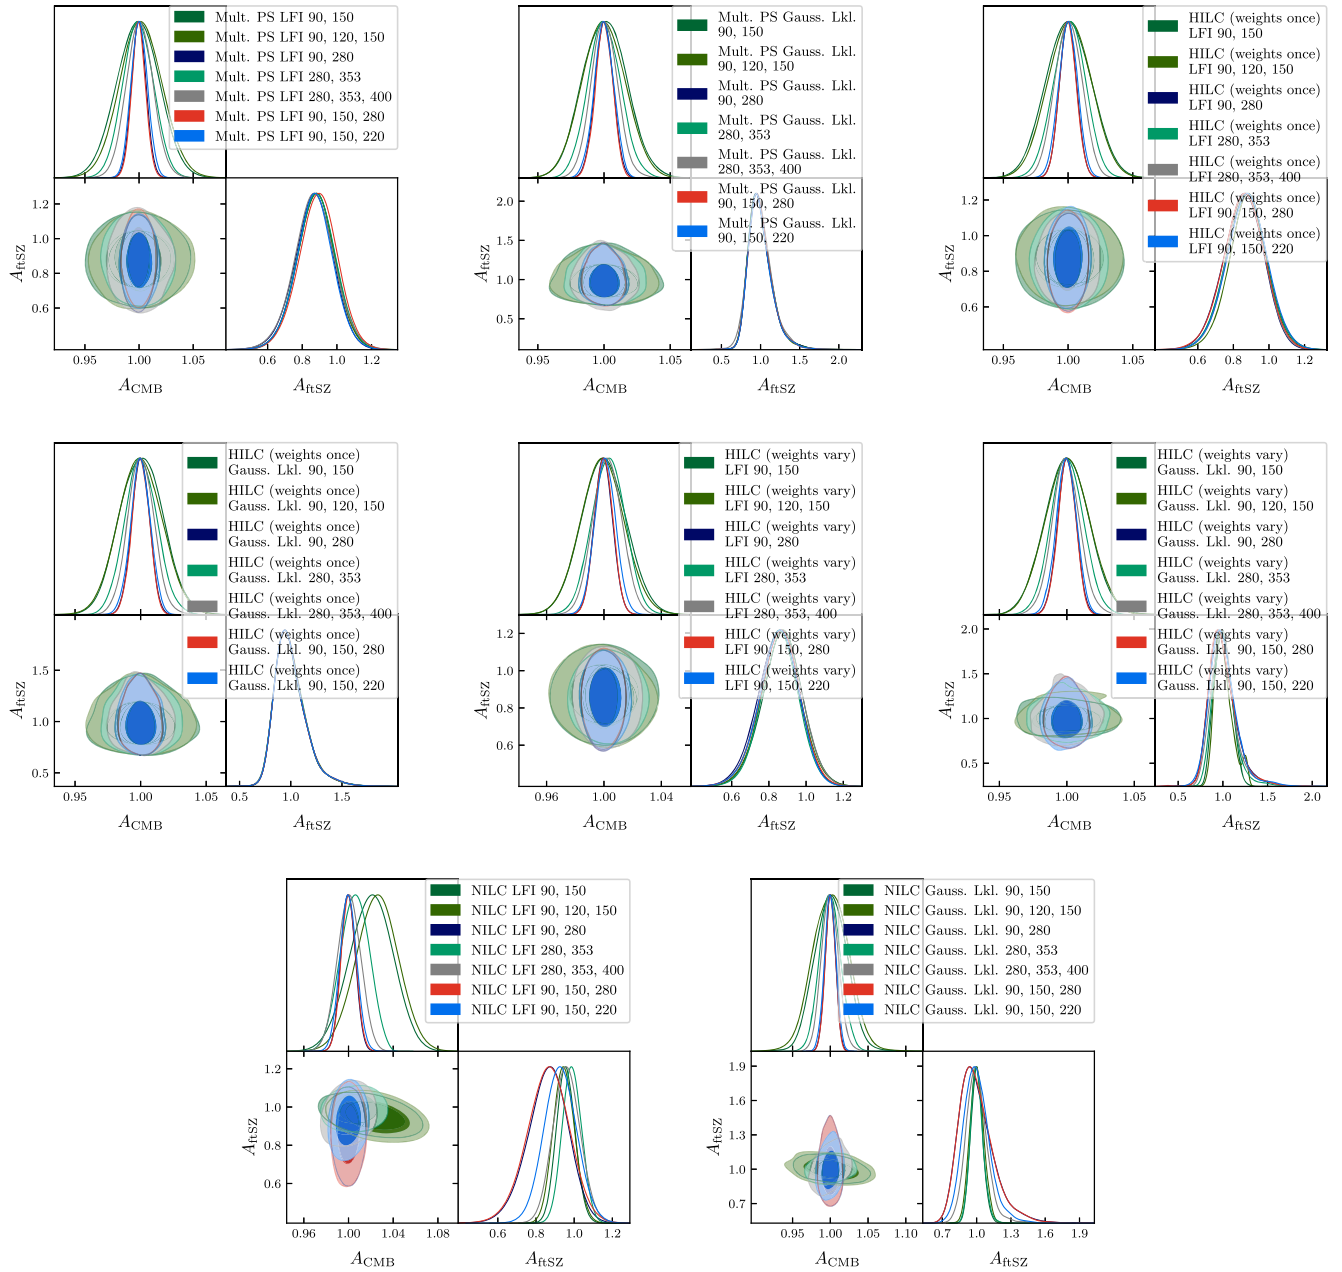


FIG. 15. Comparison of different frequency combinations used for each individual summary statistic, with both LFI and a Gaussian likelihood. For the HILC (weights once) Gaussian likelihood, we use the analytic parameter dependence.

[1] N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo, S. Basak *et al.*, *Astron. Astrophys.* **641**, A5 (2020).
 [2] J. Dunkley, E. Calabrese, J. Sievers, G. Addison, N. Battaglia, E. Battistelli, J. Bond, S. Das, M. Devlin, R. Dünner *et al.*, *J. Cosmol. Astropart. Phys.* **07** (2013) 025.

[3] S. Aiola, E. Calabrese, L. Maurin, S. Naess, B. L. Schmitt, M. H. Abitbol, G. E. Addison, P. A. R. Ade, D. Alonso, M. Amiri *et al.*, *J. Cosmol. Astropart. Phys.* **12** (2020) 047.
 [4] D. Dutcher, L. Balkenhol, P. Ade, Z. Ahmed, E. Anderes, A. Anderson, M. Archipley, J. Avva, K. Aylor, P. Barry *et al.*, *Phys. Rev. D* **104**, 022003 (2021).

- [5] P. A. R. Ade *et al.* (BICEP/Keck Collaboration), *Phys. Rev. Lett.* **127**, 151301 (2021).
- [6] Z. Atkins *et al.*, *J. Cosmol. Astropart. Phys.* **11** (2023) 073.
- [7] R. Adam *et al.* (Planck Collaboration), *Astron. Astrophys.* **594**, A8 (2016).
- [8] Z. Pan *et al.* (SPT Collaboration), *Phys. Rev. D* **108**, 122005 (2023).
- [9] K. K. Schaffer, T. M. Crawford, K. A. Aird, B. A. Benson, L. E. Bleem, J. E. Carlstrom, C. L. Chang, H. M. Cho, A. T. Crites, T. de Haan *et al.*, *Astrophys. J.* **743**, 90 (2011).
- [10] H. K. Eriksen, C. Dickinson, C. R. Lawrence, C. Baccigalupi, A. J. Banday, K. M. Górski, F. K. Hansen, P. B. Lilje, E. Pierpaoli, M. D. Seiffert *et al.*, *Astrophys. J.* **641**, 665 (2006).
- [11] H. K. Eriksen, J. B. Jewell, C. Dickinson, A. J. Banday, K. M. Górski, and C. R. Lawrence, *Astrophys. J.* **676**, 10 (2008).
- [12] C. L. Bennett, R. S. Hill, G. Hinshaw, M. R. Nolta, N. Odegard, L. Page, D. N. Spergel, J. L. Weiland, E. L. Wright, M. Halpern *et al.*, *Astrophys. J. Suppl. Ser.* **148**, 97 (2003).
- [13] M. Tegmark, A. de Oliveira-Costa, and A. J. Hamilton, *Phys. Rev. D* **68**, 123523 (2003).
- [14] H. K. Eriksen, A. J. Banday, K. M. Górski, and P. B. Lilje, *Astrophys. J.* **612**, 633 (2004).
- [15] J. Delabrouille, J. F. Cardoso, M. Le Jeune, M. Betoule, G. Fay, and F. Guilloux, *Astron. Astrophys.* **493**, 835 (2009).
- [16] W. Coulton *et al.* (ACT Collaboration), *Phys. Rev. D* **109**, 063530 (2024).
- [17] F. McCarthy and J. C. Hill, [arXiv:2307.01043](https://arxiv.org/abs/2307.01043).
- [18] J. Chandran, M. Remazeilles, and R. B. Barreiro, [arXiv:2305.10193](https://arxiv.org/abs/2305.10193).
- [19] L. E. Bleem *et al.* (SPT-SZ Collaboration), *Astrophys. J. Suppl. Ser.* **258**, 36 (2022).
- [20] M. Remazeilles, N. Aghanim, and M. Douspis, *Mon. Not. R. Astron. Soc.* **430**, 370 (2013).
- [21] K. Wolz *et al.*, *Astron. Astrophys.* **686**, A16 (2024).
- [22] A. Carones, M. Migliaccio, G. Puglisi, C. Baccigalupi, D. Marinucci, N. Vittorio, and D. Poletti, [arXiv:2212.04456](https://arxiv.org/abs/2212.04456).
- [23] K. Abazajian *et al.*, [arXiv:1907.04473](https://arxiv.org/abs/1907.04473).
- [24] Kristen M. Surrao and J. Colin Hill, preceding paper, *Phys. Rev. D* **110**, 063509 (2024).
- [25] P. Lemos, M. Cranmer, M. Abidi, C. Hahn, M. Eickenberg, E. Massara, D. Yallup, and S. Ho, *Mach. Learn. Sci. Technol.* **4**, 01LT01 (2023).
- [26] J. Alsing, T. Charnock, S. Feeney, and B. Wandelt, *Mon. Not. R. Astron. Soc.* **488**, 4440 (2019).
- [27] G.-J. Wang, C. Cheng, Y.-Z. Ma, J.-Q. Xia, A. Abebe, and A. Beesham, [arXiv:2306.11102](https://arxiv.org/abs/2306.11102).
- [28] G. Hinshaw *et al.* (WMAP Collaboration), *Astrophys. J. Suppl. Ser.* **148**, 135 (2003).
- [29] N. Aghanim *et al.* (Planck Collaboration), *Astron. Astrophys.* **641**, A6 (2020); **652**, C4(E) (2021).
- [30] S. K. Choi, M. Hasselfield, S.-P. P. Ho, B. Koopman, M. Lungu, M. H. Abitbol, G. E. Addison, P. A. R. Ade, S. Aiola, D. Alonso *et al.*, *J. Cosmol. Astropart. Phys.* **12** (2020) 045.
- [31] <https://github.com/kmsurrao/NILC-Inference-Pipeline>.
- [32] P. A. R. Ade *et al.* (Planck Collaboration), *Astron. Astrophys.* **571**, A12 (2014).
- [33] P. A. R. Ade *et al.* (Planck Collaboration), *Astron. Astrophys.* **571**, A21 (2014).
- [34] R. Adam, P. A. R. Ade, N. Aghanim, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro *et al.* (Planck Collaboration), *Astron. Astrophys.* **594**, A9 (2016).
- [35] N. Aghanim, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, J. G. Bartlett, N. Bartolo *et al.* (Planck Collaboration), *Astron. Astrophys.* **594**, A22 (2016).
- [36] F. McCarthy and J. C. Hill, [arXiv:2308.16260](https://arxiv.org/abs/2308.16260).
- [37] F. J. Narcowich, P. Petrushev, and J. D. Ward, *SIAM J. Math. Anal.* **38**, 574 (2006).
- [38] D. Marinucci, D. Pietrobon, A. Balbi, P. Baldi, P. Cabella, G. Kerkycharian, P. Natoli, D. Picard, and N. Vittorio, *Mon. Not. R. Astron. Soc.* **383**, 539 (2007).
- [39] F. Guilloux, G. Fay, and J.-F. Cardoso, [arXiv:0706.2598](https://arxiv.org/abs/0706.2598).
- [40] X. Chen and E. L. Wright, *Astrophys. J.* **694**, 222 (2009).
- [41] M. Remazeilles, J. Delabrouille, and J.-F. Cardoso, *Mon. Not. R. Astron. Soc.* **410**, 2481 (2011).
- [42] K. Cranmer, J. Brehmer, and G. Louppe, *Proc. Natl. Acad. Sci. U.S.A.* **117**, 30055 (2020).
- [43] G. Papamakarios, D. C. Sterratt, and I. Murray, [arXiv:1805.07226](https://arxiv.org/abs/1805.07226).
- [44] M. Glöckler, M. Deistler, and J. H. Macke, [arXiv:2203.04176](https://arxiv.org/abs/2203.04176).
- [45] J. Boelts, J.-M. Lueckmann, R. Gao, and J. H. Macke, *eLife* **11**, e77220 (2022).
- [46] J. Hermans, V. Begy, and G. Louppe, [arXiv:1903.04057](https://arxiv.org/abs/1903.04057).
- [47] C. Durkan, I. Murray, and G. Papamakarios, [arXiv:2002.03712](https://arxiv.org/abs/2002.03712).
- [48] A. Delaunoy, J. Hermans, F. Rozet, A. Wehenkel, and G. Louppe, [arXiv:2208.13624](https://arxiv.org/abs/2208.13624).
- [49] B. K. Miller, C. Weniger, and P. Forré, [arXiv:2210.06170](https://arxiv.org/abs/2210.06170).
- [50] G. Papamakarios and I. Murray, [arXiv:1605.06376](https://arxiv.org/abs/1605.06376).
- [51] J.-M. Lueckmann, P. J. Goncalves, G. Bassetto, K. Öcal, M. Nonnenmacher, and J. H. Macke, [arXiv:1711.01861](https://arxiv.org/abs/1711.01861).
- [52] D. S. Greenberg, M. Nonnenmacher, and J. H. Macke, [arXiv:1905.07488](https://arxiv.org/abs/1905.07488).
- [53] M. Deistler, P. J. Goncalves, and J. H. Macke, [arXiv:2210.04815](https://arxiv.org/abs/2210.04815).
- [54] G. Papamakarios, T. Pavlakou, and I. Murray, [arXiv:1705.07057](https://arxiv.org/abs/1705.07057).
- [55] B. Uria, M.-A. Côté, K. Gregor, I. Murray, and H. Larochelle, [arXiv:1605.02226](https://arxiv.org/abs/1605.02226).
- [56] D. Jimenez Rezende and S. Mohamed, [arXiv:1505.05770](https://arxiv.org/abs/1505.05770).
- [57] M. Germain, K. Gregor, I. Murray, and H. Larochelle, [arXiv:1502.03509](https://arxiv.org/abs/1502.03509).
- [58] Y. B. Zeldovich and R. A. Sunyaev, *Astrophys. Space Sci.* **4**, 301 (1969).
- [59] R. A. Sunyaev and Y. B. Zeldovich, *Astrophys. Space Sci.* **7**, 3 (1970).
- [60] G. Stein, M. A. Alvarez, J. R. Bond, A. van Engelen, and N. Battaglia, *J. Cosmol. Astropart. Phys.* **10** (2020) 012.
- [61] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann, *Astrophys. J.* **622**, 759 (2005).
- [62] A. Zonca, L. Singer, D. Lenz, M. Reinecke, C. Rosset, E. Hivon, and K. Gorski, *J. Open Source Software* **4**, 1298 (2019).

- [63] J. L. Tinker, A. V. Kravtsov, A. Klypin, K. Abazajian, M. S. Warren, G. Yepes, S. Gottlober, and D. E. Holz, *Astrophys. J.* **688**, 709 (2008).
- [64] N. Battaglia, J. R. Bond, C. Pfrommer, and J. L. Sievers, *Astrophys. J.* **758**, 75 (2012).
- [65] L. Knox, *Phys. Rev. D* **52**, 4307 (1995).
- [66] J. Hartlap, P. Simon, and P. Schneider, *Astron. Astrophys.* **464**, 399 (2007).
- [67] J. A. Nelder and R. Mead, *Comput. J.* **7**, 308 (1965).
- [68] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, *Publ. Astron. Soc. Pac.* **125**, 306 (2013).
- [69] A. Tejero-Cantero, J. Boelts, M. Deistler, J.-M. Lueckmann, C. Durkan, P. J. Gonçalves, D. S. Greenberg, and J. H. Macke, *J. Open Source Software* **5**, 2505 (2020).
- [70] L. Biewald, Experiment tracking with weights and biases, software available from <https://www.wandb.com/> (2020).
- [71] M. Cranmer, [arXiv:2305.01582](https://arxiv.org/abs/2305.01582).
- [72] T. Louis, E. Grace, M. Hasselfield, M. Lungu, L. Maurin, G. E. Addison, P. A. R. Ade, S. Aiola, R. Allison, M. Amiri *et al.*, *J. Cosmol. Astropart. Phys.* **06** (2017) 031.
- [73] J. A. Kable, G. E. Addison, and C. L. Bennett, *Astrophys. J.* **888**, 26 (2020).
- [74] M. von Wietersheim-Kramsta, K. Lin, N. Tessore, B. Joachimi, A. Loureiro, R. Reischke, and A. H. Wright, [arXiv:2404.15402](https://arxiv.org/abs/2404.15402).
- [75] M. S. Madhavacheril, J. C. Hill, S. Næss, G. E. Addison, S. Aiola, T. Baildon, N. Battaglia, R. Bean, J. R. Bond, E. Calabrese *et al.*, *Phys. Rev. D* **102**, 023534 (2020).
- [76] Y. Akrami *et al.* (Planck Collaboration), *Astron. Astrophys.* **641**, A11 (2020).
- [77] N. Krachmalnicoff and G. Puglisi, *Astrophys. J.* **911**, 42 (2021).
- [78] B. Thorne, L. Knox, and K. Prabhu, *Mon. Not. R. Astron. Soc.* **504**, 2603 (2021).
- [79] B. Regaldo-Saint Blancard, F. Levrier, E. Allys, E. Bellomi, and F. Boulanger, *Astron. Astrophys.* **642**, A217 (2020).
- [80] A. Lewis, [arXiv:1910.13970](https://arxiv.org/abs/1910.13970).
- [81] C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith *et al.*, *Nature (London)* **585**, 357 (2020).
- [82] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright *et al.*, *Nat. Methods* **17**, 261 (2020).
- [83] J. D. Hunter, *Comput. Sci. Eng.* **9**, 90 (2007).
- [84] T. P. Robitaille, E. J. Tollerud, P. Greenfield, M. Droettboom, E. Bray, T. Aldcroft, M. Davis, A. Ginsburg, A. M. Price-Whelan *et al.* (Astropy Collaboration), *Astron. Astrophys.* **558**, A33 (2013).
- [85] A. M. Price-Whelan, B. M. Sipőcz, H. M. Günther, P. L. Lim, S. M. Crawford, S. Conseil, D. L. Shupe, M. W. Craig, N. Dencheva *et al.* (Astropy Collaboration), *Astron. J.* **156**, 123 (2018).
- [86] A. M. Price-Whelan, P. L. Lim, N. Earl, N. Starkman, L. Bradley, D. L. Shupe, A. A. Patil, L. Corrales, C. E. Brasseur *et al.* (Astropy Collaboration), *Astrophys. J.* **935**, 167 (2022).
- [87] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga *et al.*, [arXiv:1912.01703](https://arxiv.org/abs/1912.01703).
- [88] J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah, *SIAM Rev.* **59**, 65 (2017).
- [89] S. Hamimeche and A. Lewis, *Phys. Rev. D* **77**, 103013 (2008).
- [90] D. Coe, [arXiv:0906.4123](https://arxiv.org/abs/0906.4123).