

Constraining cosmology with thermal Sunyaev-Zel'dovich maps: Minkowski functionals, peaks, minima, and moments

Alina Sabyr^{1,*}, J. Colin Hill², and Zoltán Haiman^{1,2}

¹*Department of Astronomy, Columbia University, New York, New York 10027, USA*

²*Department of Physics, Columbia University, New York, New York 10027, USA*



(Received 29 October 2024; accepted 15 April 2025; published 28 May 2025)

The thermal Sunyaev-Zel'dovich effect (tSZ) is a sensitive probe of cosmology, as it traces the abundance of galaxy clusters and groups in the late-time Universe. Upcoming cosmic microwave background experiments such as the Simons Observatory (SO) and CMB-S4 will provide low-noise and high-resolution component-separated tSZ maps covering a large sky fraction. The tSZ signal is highly non-Gaussian; therefore, higher-order statistics are needed to optimally extract information from these maps. In this work, we study the cosmological constraining power of several tSZ statistics—Minkowski functionals (MFs), peaks, minima, and moments—that have yielded promising results in capturing non-Gaussian information from other cosmological data. Using a large suite of halo-model-based tSZ simulations with varying Ω_c and σ_8 (154 cosmologies and over 800,000 maps, each $10.5 \times 10.5 \text{ deg}^2$), we show that by combining these observables, we can achieve $\approx 29\times$ tighter constraints compared to using the tSZ power spectrum alone in an idealized noiseless case, with the MFs dominating the constraints. We show that much of the MF constraining power arises from halos below the detection threshold of cluster surveys, suggesting promising synergies with cluster-count analyses. Finally, we demonstrate that these statistics have the potential to deliver tight constraints even in the presence of noise. For example, using post-component-separation tSZ noise expected for SO, we obtain $\approx 1.6\times$ and $\approx 1.8\times$ tighter constraints than the power spectrum with MFs and all statistics combined, respectively. We show that the constraints from MFs approach the noiseless case for white-noise levels $\lesssim 1 \mu\text{K-arcmin}$.

DOI: 10.1103/PhysRevD.111.103536

I. INTRODUCTION

Cosmic microwave background (CMB) primary anisotropies have allowed us to place extraordinarily tight constraints on the parameters characterizing the Λ CDM model (e.g., *WMAP* [1], *Planck* [2]). Current and next-generation high-resolution multifrequency CMB anisotropy experiments such as the Atacama Cosmology Telescope (ACT),¹ South Pole Telescope (SPT),² Simons Observatory (SO),³ and CMB-S4,⁴ are now also able to measure the secondary anisotropies—those arising from processes after recombination—to high precision.

One of these secondary effects, which is among the leading contributions to the CMB temperature fluctuations on arcminute scales, is the thermal Sunyaev-Zel'dovich (tSZ) effect—a spectral distortion in the CMB produced via the inverse-Compton scattering of CMB photons off free, energetic electrons, primarily located in galaxy groups and

clusters [3,4]. Under the assumption of self-similarity, the observed tSZ effect scales with dark matter halo mass as $\propto M^{5/3}$ [5].

The unique spectral signature of the tSZ effect and its sensitivity to the abundance of massive halos—independent of redshift—make it a powerful cosmological probe. It can be used to identify clusters in blind surveys, providing nearly mass-limited samples, which can then be used to constrain cosmological parameters via the halo mass function (e.g., [6–11]). The signal can also be studied statistically across the sky through CMB temperature maps at a given frequency or component-separated tSZ maps via multifrequency CMB observations (e.g., [12–22]). In this paper, we will focus on the latter, statistical approach.

So far, the tSZ power spectrum (e.g., [12,15,19–23]), bispectrum (e.g., [13,23]), one-point probability distribution (PDF) (e.g., [18,23,24]), moments (up to fourth order, e.g., [14,16]) and skew-spectra (e.g., [25,26]) have been studied. With respect to the cosmological parameters, the tSZ signal has been shown to have high sensitivity to the amplitude of linear matter density fluctuations at $z = 0$, σ_8 , through the steep dependence of the strength of the signal on halo mass, and therefore strong response to the halos in the exponential tail of the mass function (e.g., [12,13,16]).

*Contact author: as6131@columbia.edu

¹<https://act.princeton.edu/>

²<https://pole.uchicago.edu/public/Home.html>

³<https://simonsobservatory.org/>

⁴<https://cmb-s4.org/>




FIG. 1. Comparison of a simulated halo-model-based tSZ map (at 0.1 arcmin resolution) to a Gaussian random field generated from its power spectrum. Notice the “hot spots” on the left, missing in the Gaussian case.

The tSZ signal is also sensitive to the dark energy equation of state parameter, w , massive neutrinos, M_ν , and primordial non-Gaussianity through their effects on the growth of large-scale structure (e.g., [15,22,27]).

The tSZ field is highly non-Gaussian since it is sourced by massive structures in the late-time Universe (see Fig. 1). This motivates the use of summary statistics beyond the power spectrum. Indeed, it has been shown that higher-point tSZ statistics scale more steeply with cosmological parameters [13,16] and that the one-point PDF provides tighter constraints than the skewness, for example [14,18].

The primary goal of this work is to perform a systematic study of several summary statistics as a step toward: (1) finding an optimal tSZ statistic and (2) characterizing the full information content of the tSZ field. In this paper, we focus on two cosmological parameters: the cold dark matter density, Ω_c , and σ_8 . In particular, we compare the cosmological constraints achieved using the power spectrum to those obtained from four sets of observables: Minkowski functionals, peaks, minima, and moments. To our knowledge, this is the first time that most of these statistics have been applied to tSZ maps. These exact descriptors have been shown to effectively extract non-Gaussian information from other cosmological fields such as, for example, weak gravitational lensing convergence maps (e.g., [28–41]).⁵ We note that the tSZ effect depends on the details of the gas physics and, in practice, parameters associated with the intracluster medium (ICM) physics would need to be simultaneously fit along with cosmological parameters. However, we focus on cosmology here and leave the exploration of the uncertainty and sensitivity related to astrophysical parameters to follow-up work.

⁵Some of these summary statistics have also been applied to galaxy clustering (e.g., [42]) and have been used to place limits on primordial non-Gaussianity from primary anisotropies (e.g., [43,44]).

The steep scaling of the tSZ effect with halo mass poses both a computational challenge and an advantage. In order to properly capture the sample variance and the mean for each of these summary statistics, a large number of realizations and simulation size are needed. On the other hand, this also means that the signal is dominated by the 1-halo or the Poisson term rather than the clustering of halos [45]. For example, it has been shown for the power spectrum that the clustering effects are negligible for $\ell > 300$ [15] and are also small for the one-point PDF [24]. The effect is expected to be even less important for higher-point statistics, which scale more steeply with mass. Therefore, it is not as crucial to simulate realistic clustering of halos unlike for other cosmological fields such as, for example, weak gravitational lensing [46,47], and simplified simulation approaches can be sufficiently accurate.

Since it is time-prohibitive to perform an exhaustive study using a large suite of hydrodynamical simulations and the Poisson term dominates the tSZ signal, we make use of “simplified simulations” in this work. Namely, we simulate the tSZ maps by Poisson sampling halos from a given mass function, assigning a y -profile, and placing halos randomly in a given map [18,24]. Using this fast approach, we are able to generate a large suite of simulations with sufficient size and number of random realizations to reach convergence across all the observables and to properly estimate the covariance matrices in this work. Using this suite of simulations, we study the constraining power of each of the five summary statistics both in noiseless and noisy cases. For the latter, we use post-component-separation tSZ noise power spectra for upcoming CMB experiments.

Our main findings are summarized in Fig. 7 (noiseless, Table III) and Fig. 14 (noisy, Table VIII). In the noiseless case, we find that combining all five summary statistics gives $\approx 29\times$ tighter constraints than using the power spectrum alone, with MFs significantly outperforming all

other descriptors. The power spectrum + MFs and power spectrum + peaks yield $23\times$ and $3.4\times$ tighter constraints than the power spectrum alone, respectively. In the noisy case, we can achieve $\approx 1.8\times$ tighter constraints with all the descriptors, as compared to those from the power spectrum alone. MFs still perform better than any other summary statistic: the power spectrum + MFs achieve $\approx 1.7\times$ improvement over the power spectrum alone, respectively.

The remainder of this paper is organized as follows. We describe the details of the simulation suite in Sec. II, provide an overview of the statistical descriptors in Sec. III, and outline the inference methods in Sec. IV. We present the results on smoothed noiseless maps in Sec. V and the results on noisy maps in Sec. VI. We summarize our findings and conclude in Sec. VII.

We assume a flat Λ CDM cosmology throughout this paper with parameters set to be consistent with *Planck* [2], except that we assume no massive neutrinos: $h = 0.674$, where $h \equiv H_0/(100 \text{ km s}^{-1} \text{ Mpc}^{-1})$ and H_0 is the Hubble constant; baryon density $\Omega_b = 0.0493$; spectral index $n_s = 0.965$, and optical depth to reionization $\tau_{\text{reio}} = 0.054$. We set our fiducial cosmology to have $\Omega_c = 0.264$, and $\sigma_8 = 0.811$.

II. SIMULATIONS

In the nonrelativistic limit, the CMB temperature fluctuations generated due to inverse-Compton scattering at a given sky position $\hat{\mathbf{n}}$ can be written as [3,4]

$$\frac{\Delta T(\nu, \hat{\mathbf{n}})}{T_{\text{CMB}}} = g(\nu)y(\hat{\mathbf{n}}), \quad \text{where } g(\nu) = x \coth\left(\frac{x}{2}\right) - 4. \quad (1)$$

Here, ν is the photon frequency, $g(\nu)$ is the tSZ spectral function, $x = (h\nu)/(k_B T_{\text{CMB}})$, h is Planck's constant, k_B is Boltzmann's constant, and T_{CMB} is the CMB temperature today. The strength of the signal is parameterized by the Compton- y parameter, which is proportional to the integral of the electron pressure along the line of sight (LOS)

$$\begin{aligned} y(\hat{\mathbf{n}}) &= \frac{\sigma_T}{m_e c^2} \int n_e(\hat{\mathbf{n}}, l) k_B T_e(\hat{\mathbf{n}}, l) dl \\ &= \frac{\sigma_T}{m_e c^2} \int P_e(\hat{\mathbf{n}}, l) dl \end{aligned} \quad (2)$$

where n_e , T_e , P_e , and m_e are the electron number density, temperature, pressure, and mass, respectively; σ_T is the Thomson cross section, and c is the speed of light. The simulated maps that we use in this work are directly in units of y and we therefore denote various quantities in this work in terms of y .

A. tSZ maps

To simulate the tSZ signal, we use the publicly available code HMPDF [24,46].⁶ The maps are generated by Poisson sampling halos from the Tinker *et al.* [48] halo mass function in discrete bins of mass and redshift. The halo mass function is computed using the linear matter power spectrum from CLASS [49]. Each halo is assigned a Compton- y profile using the pressure profile fitting function from Battaglia *et al.* [50]. Both halo mass and the pressure profile fitting functions are defined for a spherical overdensity mass, M_Δ , corresponding to an enclosed mass within r_Δ with interior density equal to Δ times the critical (c) or mean matter density (m) at that redshift. To convert between $M_{200\text{m}}$ and $M_{200\text{c}}$, the concentration-mass relation from Duffy *et al.* [51] is used. Unless otherwise noted, we set the following bounds for the mass function: $10^{11} M_\odot \leq M \leq 10^{16} M_\odot$, $0.005 \leq z \leq 6$. For the y -profiles, we apply a radial cutoff at $r_{\text{out}} = 2r_{\text{vir}}$, where r_{vir} is the virial radius. We use 150 bins in both redshift and mass (i.e., $N_z = N_M = 150$), with bins spaced according to Gauss-Legendre quadrature within $(z_{\text{min}}, z_{\text{max}})$ and $(\ln M_{\text{min}}, \ln M_{\text{max}})$ to populate our maps, based on the convergence of the mean value of each of the summary statistics applied in this work (see Appendix A for a description of the simulation convergence tests).

For computational speed, we simulate random realizations of maps that cover 10% of the sky at 0.1 arcmin resolution and cut each large map into 36 smaller ones. To remove any possible correlations at the edges, we trim 50 pixels on each side of the small maps. The final size of each map is $10.535 \times 10.535 \text{ deg}^2$ (6321×6321 pixels²). We apodize each small map with a cosine function, using 5 pixels for the width of the apodization at each edge, and then smooth the map using a Gaussian kernel with $\theta_{\text{FWHM}} = 1.4$ arcmin. An example simulated, smoothed map is shown in panel (a) of Fig. 4.

We choose the pixel resolution to be 0.1 arcmin based on a comparison of the high- ℓ tSZ power spectra, C_ℓ^{yy} , between maps generated at several different resolutions and smoothed with the same Gaussian kernel. To validate our simulations, we compare the power spectrum estimated from the maps to that computed from the halo-model code CLASS_SZ [12,22,27,52,53]⁷ at the fiducial cosmology (see Appendix A).

Since in this work we only use the power spectra estimated from the simulated maps and, therefore, are only sensitive to the relative difference between different sets of simulations, we do not divide out the pixel window function or account for the apodization mask in the forecasts. Our map pipeline, therefore, only includes apodization and smoothing operations. We also note that we always subtract the mean value from the final maps over

⁶<https://github.com/leanderthiele/hmpdf>

⁷https://github.com/CLASS-SZ/class_sz




FIG. 2. The simulation suite used in this work consists of 154 cosmologies with varying Ω_c and σ_8 . The points are sampled more densely near the fiducial cosmology. We have $N_{\text{fid}} = 34,560$ random realizations at the fiducial cosmology (red diamond marker) and $N_{\text{sims}} = 5,112$ at every other cosmology (gray points).

which we measure each of the statistics; thus, the values in the maps are both negative and positive.

B. Simulation suite

Our simulation suite consists of 154 different cosmologies, as shown in Fig. 2, with varying parameters Ω_c and σ_8 , with $N_{\text{fid}} = 34,560$ random realizations at the fiducial cosmology and $N_{\text{sims}} = 5,112$ at all others. We set our fiducial cosmology to be at $\Omega_c = 0.264$, $\sigma_8 = 0.811$ and fix the rest of the cosmological parameters to results from *Planck* [2], but with no massive neutrinos: $h = 0.674$, $\Omega_b = 0.0493$, $n_s = 0.965$, $\tau_{\text{reio}} = 0.054$.

We choose the number of simulations at the fiducial cosmology based on an estimate of how many realizations will be needed to accurately compute the covariance for a summary statistic with dimension $d \sim \text{few hundred}$ (i.e., $N_{\text{fid}} \gtrsim d^2/2$). In Sec. V C, we show that the constraints are stable with $N_{\text{fid}} < 34,560$. We pick $N_{\text{sims}} = 5,112$ for other cosmologies based on the stability of the mean value of each summary statistic considered in this work. As described in Appendix A, the mean converges across the mass function discretization settings in HMPDF at $N_{\text{sims}} < 5,112$.

We create the suite by first running 102 cosmologies within the broad ranges $\Omega_c \in \{0.21, 0.32\}$ and $\sigma_8 \in \{0.6, 1\}$ (approximately 20% or 30 σ from the *Planck* central values, respectively) with 99 of those sampled according to the Latin hypercube (LH). After computing the initial constraints from all the statistics, we add extra cosmologies in the region near the maximum likelihood to ensure that we have a sufficient number of cosmologies to interpolate each of the summary statistics (see Sec. IV) and obtain reliable

contours. In particular, we add 20 cosmologies sampled along a line fit to the $\Omega_c - \sigma_8$ degeneracy found for the power spectrum and 20 additional points around the constraints from all the statistics combined. The latter we do by following a similar approach as in Ref. [54]. We first compute the approximate covariance using the parameter values that lie within the 1σ constraints and the confidence ellipse parameters a , b , γ following Ref. [55]:

$$\begin{aligned} a^2 &= \frac{\sigma_x^2 + \sigma_y^2}{2} + \sqrt{\frac{(\sigma_x^2 - \sigma_y^2)^2}{4} + \sigma_{xy}^2} \\ b^2 &= \frac{\sigma_x^2 + \sigma_y^2}{2} - \sqrt{\frac{(\sigma_x^2 - \sigma_y^2)^2}{4} + \sigma_{xy}^2} \\ \tan(2\gamma) &= \frac{2\sigma_{xy}}{\sigma_x^2 - \sigma_y^2} \end{aligned} \quad (3)$$

We then sample values for $r \in \{0, 1\}$, $\phi \in \{0, 2\pi\}$ and convert to Ω_c , σ_8 via

$$\begin{aligned} x &= \alpha a \cos(\phi) r^n \\ y &= \alpha b \sin(\phi) r^n \\ \Omega_c &= x \cos(\gamma) - y \sin(\gamma) + \Omega_{c,\text{fid}} \\ \sigma_8 &= x \sin(\gamma) + y \cos(\gamma) + \sigma_{8,\text{fid}} \end{aligned} \quad (4)$$

where we set $n = 1$, $\alpha = 3.5$ ($\alpha = 3.44$ corresponds to 3σ). Choosing $n > 1/2$ means that the density of samples increases toward the center. In Appendix C, we show that with these cosmologies we are able to construct an accurate interpolator for all the summary statistics and achieve stable forecasts.

Finally, we add 12 more cosmologies that can be used for Fisher forecasts, where one of the two parameters is perturbed from its fiducial value, while the other is fixed. For these we use $\delta\theta = \pm 0.5, 1, 3\%$, which correspond to $\delta\Omega_c = \pm\{0.00132, 0.00264, 0.00792\}$, $\delta\sigma_8 = \pm\{0.00406, 0.00811, 0.02433\}$.

C. Noise

To study the constraining power of the tSZ statistics in the presence of realistic noise, we generate Gaussian random fields (GRFs) using different noise power spectra, N_ℓ^{yy} . Note that here N_ℓ^{yy} refers to the nonbeam-deconvolved noise power spectra, as we always use smoothed tSZ maps to compute the forecasts.

First, we consider several experimental setups using the post-component-separation Compton- y noise power spectra forecasted for SO (baseline and goal noise levels)⁸ [56]

⁸https://github.com/simonsobs/so_noise_models/tree/master/LAT_comp_sep_noise/v3.1.0

and CMB-S4 [57],⁹ both computed assuming combination with *Planck* data. These power spectra were generated for an assumed model, which includes both the instrumental noise and foreground components. Thus, these noise power spectra include the expected residuals from the other sky components during the component-separation process to obtain tSZ maps from multifrequency CMB observations. However, note that the noise power spectra (including the residual foreground fields) are Gaussian and are not correlated with the tSZ field, complexities that we leave to future work. In all cases, we use the noise computed via the standard internal linear combination (ILC) method and assume $f_{\text{sky}} = 0.4$ (for more details, see Refs. [56,57]).

The largest-angular-scale noise is not well-captured by the ILC method due to the small number of harmonic modes at low ℓ . To avoid any ill behavior in the maps due to extrapolation of the noise power spectra at low- and high- ℓ ranges, we filter out the scales $\ell \leq 80$ and $\ell \geq 7950$ (smoothly tapering with a tanh function between $\ell = 80$ – 90 and $\ell = 7940$ – 7950) from both the tSZ maps and the noise maps (i.e., near the boundaries of the SO noise power spectra). The noise curves and their comparison to the tSZ power spectrum at the fiducial cosmology are shown in Fig. 3.

Figure 4 shows visually the impact of adding noise to the tSZ maps for the SO goal case. Panel (a) is the noiseless map described in the previous section smoothed with a $\theta_{\text{FWHM}} = 1.4$ arcmin Gaussian kernel. Panel (b) shows the effects on the tSZ map of filtering out the low- and high- ℓ modes. Panel (c) is the total map, comprised of the smoothed tSZ map plus the random noise realization. While some structure from the tSZ field still remains, the noisy map is much more similar to a GRF (see, e.g., Fig. 1).

To mitigate the effects of noise in the real-space statistics that we study, we apply a Wiener filter to the noisy tSZ maps. We compute the ℓ -space filter using the tSZ power spectrum measured at the fiducial cosmology (averaged over $N_{\text{fid}} = 34,560$ realizations) from the noiseless maps and N_{ℓ}^{yy} using $80 < \ell < 7950$:

$$F_{\ell}^{\text{Wiener}} = \frac{C_{\ell}^{yy}}{C_{\ell}^{yy} + N_{\ell}^{yy}}. \quad (5)$$

The filter preferentially upweights the modes that contain the largest tSZ signal-to-noise ratio. Since we construct the filter using the power spectrum, it is not guaranteed that it is optimal for the real-space non-Gaussian statistics that we analyze. Nonetheless, it still suppresses the modes that are clearly dominated by noise.

⁹See Fig. 20 of Ref. [57].




FIG. 3. Top: post-component-separation tSZ noise for SO baseline (goal) noise levels in gray (black) and CMB-S4 (cyan) compared to the tSZ power spectrum (blue) computed at the fiducial cosmology. Combination with *Planck* data is assumed in these noise power spectra and the increase in noise around $\ell \approx 1000$ – 1500 shows the transition between the *Planck*- and SO/CMB-S4-dominated regimes (i.e., low- ℓ modes are noisy for SO/CMB-S4 due to atmospheric effects). White-noise power spectra for different noise levels are plotted with dashed curves in orange. Note that the signal curve here includes smoothing with a FWHM = 1.4 arcmin beam. Bottom: Wiener filter computed for the SO goal noise.

We plot the filter for SO goal noise levels in the bottom panel of Fig. 3 and show its effects on the simulated maps in the fourth panel of Fig. 4. We normalize the filter via dividing it by its maximum value.

In addition to noise power spectra from realistic experimental setups, we also consider different levels of white noise: $\Delta T = 0.1, 1$, and $10 \mu\text{K-arcmin}$. To generate GRFs, we convert from CMB temperature units, $\mu\text{K-arcmin}$, to Compton-y units using the tSZ spectral function $g(\nu)$ in Eq. (1) and $T_{\text{CMB}} = 2.7255$ K. We convert the units at $\nu = 150$ GHz. The white noise maps are then generated assuming a flat power spectrum with an amplitude equal to the noise level in Compton-y units $\times$ radians squared. As in the case of the postcomponent noise power spectra, we apply the same low- and high- ℓ filter and a Wiener filter computed for each of the flat noise power spectra cases. We show the white noise power spectra in Fig. 3 and the combined tSZ plus white noise maps in Fig. 20. Note that for white noise levels of 0.1 or $1 \mu\text{K-arcmin}$, the maps closely resemble the noiseless case, as expected from the signal versus noise power spectra comparison shown in Fig. 3.




FIG. 4. Simulated maps used to obtain various forecasts in this work at different steps of processing. (a) Noiseless tSZ map smoothed with a Gaussian kernel with $\theta_{\text{FWHM}} = 1.4$ arcmin ($t\text{SZ map} + \theta_{\text{FWHM}} = 1.4$ arcmin smoothing), (b) low- and high- ℓ modes filtered out ($+\ell \leq 80, \ell \geq 7950$ filtered out), (c) adding a Gaussian random field generated using the post-component-separation tSZ noise power spectra forecasted for the SO goal noise levels (+ noise (SO goal)), and (d) effect of applying a Wiener filter to the noisy map (+ wiener filter).

TABLE I. We list the fiducial binning choices and dimensions for each of the summary statistics both for the noiseless and noisy cases (i.e., for the results shown in Figs. 7 and 14), where $\sigma_{\text{fid}} = 1.5 \times 10^{-6}$ and $\sigma_{\text{fid}}^{\text{SO,goal}} = 1.9 \times 10^{-6}$ correspond to the average rms values computed at the fiducial cosmology.

Observable	Noiseless		Noisy	
	Dimension	Bin edges	Dimension	Bin edges
Power spectrum	77	Linearly spaced, $\ell \in \{125, 7825\}$	77	Linearly spaced, $\ell \in \{125, 7825\}$
MFs	3×20	Linearly spaced, $\nu \in \{-1\sigma_{\text{fid}}, 3\sigma_{\text{fid}}\}$	3×30	Linearly spaced, $\nu \in \{-3\sigma_{\text{fid}}^{\text{SO,goal}}, 3\sigma_{\text{fid}}^{\text{SO,goal}}\}$
Peaks	9	$[-100, -0.3, -0.2, -0.1, 0, 0.2, 0.3, 0.6, 1, 1000]\sigma_{\text{fid}}$	9	$[-100, -0.3, -0.2, -0.1, 0, 0.2, 0.3, 0.6, 1, 1000]\sigma_{\text{fid}}^{\text{SO,goal}}$
Minima	7	$[-100, -0.57, -0.5, -0.47, -0.42, -0.36, -0.26, 100]\sigma_{\text{fid}}$	7	$[-100, -0.57, -0.5, -0.47, -0.42, -0.36, -0.26, 100]\sigma_{\text{fid}}^{\text{SO,goal}}$
Moments	4	...	4	...

III. SUMMARY STATISTICS

In this paper, we employ five different summary statistics: the angular power spectrum, Minkowski functionals, peaks, minima, and moments. We use the implementation in LensTools [58]¹⁰ and describe each of them and our fiducial choices in this section. We note that some of the statistics are measured for a set of bins. To define the bins, we generally use values defined in terms of the average rms of the tSZ maps at the fiducial cosmology. These values are $\sigma_{\text{fid}} = 1.5 \times 10^{-6}$ and $\sigma_{\text{fid}}^{\text{SO,goal}} = 1.9 \times 10^{-6}$, for the fiducial noiseless and noisy cases, respectively. These values correspond to the average rms computed over 5,184 maps at the fiducial cosmology.¹¹ We summarize the fiducial choices and the size of each of the summary statistic data vectors in Table I.

¹⁰<https://github.com/apetri/LensTools>

¹¹The rms values are computed over 5,184 map realizations based on the number of simulations used for the initial convergence tests prior to running the suite.

A. Angular power spectrum

We use the tSZ power spectrum as a baseline statistic with which to compare the constraining power of each of the summary statistics. It is defined as

$$\langle \tilde{y}_{\ell}^* \tilde{y}_{\ell'} \rangle = (2\pi)^2 \delta_D(\ell - \ell') C_{\ell}^{yy} \quad (6)$$

where ℓ is the 2D Fourier wave vector, $\tilde{y}$ indicates the Fourier transform of the field, and δ_D is the Dirac delta function. We average the pixel values in equal-size linearly spaced $\ell = |\ell|$ bins between $\ell \in \{25, 7925\}$ (bin width equal to $\Delta\ell = 100$). We remove the lowest and highest ℓ bins to avoid any aliasing effects from applying an apodization mask to our maps and end up with a total of 77 bins with the first and the last bins centered at $\ell_{\text{min}} = 175$ and $\ell_{\text{max}} = 7775$.

B. Minkowski functionals

For a D-dimensional field, one can compute $D + 1$ Minkowski functionals (MFs), which fully characterize

the global morphological properties (i.e., invariant to rotation, translation, and additive) of that field [59–61].¹² Thus, for a 2D tSZ map, we can measure three MFs: $V_0(\nu) \sim \text{area}$, $V_1(\nu) \sim \text{contour length}$, and $V_2(\nu) \sim \text{Euler characteristic}$, as a function of some threshold value ν . More precisely,

$$V_0(\nu) = \frac{1}{A} \int_{\Sigma(\nu)} da, \quad (7)$$

where A is the total area of the map or the total number of pixels, da is the area element, and $\Sigma(\nu) = \{y > \nu\}$ defines the excursion set over which the MF is measured. V_0 is, thus, the one-point cumulative distribution function (one-point CDF). Next, $V_1(\nu)$ is defined as

$$V_1(\nu) = \frac{1}{4A} \int_{\partial\Sigma(\nu)} dl, \quad (8)$$

where dl is the boundary length element and $\partial\Sigma(\nu)$ is the excursion set boundary. Finally, $V_2(\nu)$ is defined as the integral of the curvature of the boundary $\mathcal{K}$:

$$V_2(\nu) = \frac{1}{2\pi A} \int_{\partial\Sigma(\nu)} \mathcal{K} dl \quad (9)$$

In flat space, $V_2(\nu)$ is equivalent to the Euler characteristic [63] and for a closed surface, it directly relates to the genus.

In practice, MFs can be efficiently computed as follows:

$$\begin{aligned} V_0(\nu) &= \frac{1}{A} \int_A \Theta(y(\mathbf{p}) - \nu) di dj \\ V_1(\nu) &= \frac{1}{4A} \int_A \delta_D(y(\mathbf{p}) - \nu) \sqrt{y_i^2 + y_j^2} di dj \\ V_2(\nu) &= \frac{1}{2\pi A} \int_A \delta_D(y(\mathbf{p}) - \nu) \\ &\quad \times \frac{2y_i y_j y_{ij} - y_i^2 y_{jj} - y_j^2 y_{ii}}{y_i^2 + y_j^2} di dj \end{aligned} \quad (10)$$

where $\mathbf{p}$ denotes the coordinate position (i, j) , Θ is the Heaviside step function, and $y_{i/j}$ is the gradient of the field with respect to coordinates i/j (i.e., y_{ii} corresponds to the second derivative). Thus, what is required in practice are simply the first and second derivatives of the field. In the cases where the gradients vanish (i.e., the denominator in V_2), the integrand is replaced by $(2y_{ij} - y_{jj} - y_{ii})/2$. The δ_D -functions are approximated by dividing out $d\nu$ (see more details in Ref. [30]). We check the estimators by

computing MFs on a set of GRFs for which an analytic form is known [29,63–65].

In our fiducial setup on the noiseless maps, we measure each of the MFs using 20 linearly spaced bins between $\nu \in \{\nu_{\min}, \nu_{\max}\} \in \{-1\sigma_{\text{fid}}, 3\sigma_{\text{fid}}\}$ (i.e., the MF data vector contains 60 elements). For the fiducial noisy case, we use $\nu \in \{\nu_{\min}, \nu_{\max}\} \in \{-3\sigma_{\text{fid}}^{\text{SO,goal}}, 3\sigma_{\text{fid}}^{\text{SO,goal}}\}$ and 30 bins for each of the MFs. We choose these ranges to fully capture the shape for each of the MFs. In Appendix C, we show that increasing the number of bins does not impact the forecasts from MFs.

C. Peaks and minima

Peaks (minima) are defined as pixels in the map that are higher (lower) than their 8 neighboring pixels. We locate the peaks and minima in the maps and bin them by their values. Our observable is then a peak (minima) count histogram. We use 9 and 7 bins for the peaks and minima, respectively, with edges defined as follows:

$$\begin{aligned} x_b^{\text{peaks}} &= [-100, -0.3, -0.2, -0.1, 0, 0.2, 0.3, 0.6, 1., 1000]\sigma_0 \\ x_b^{\text{minima}} &= [-100, -0.57, -0.5, -0.47, -0.42, \\ &\quad -0.36, -0.26, 100]\sigma_0 \end{aligned} \quad (11)$$

where we use $\sigma_0 = \sigma_{\text{fid}}$ and $\sigma_0 = \sigma_{\text{fid}}^{\text{SO,goal}}$ for the noiseless and noisy cases, respectively. To choose these exact bins, we take a realization of a fiducial map and bin the peaks (minima) into 10 bins in such a way that each bin has the same number of peak (minima) counts. We adjust the bin edges to ensure the stability of the mean value of each bin with respect to the numerical settings for the maps and the number of simulations (see Appendix A for more details about the convergence). Although the number of peaks (minima) is expected to vary significantly across different cosmologies, we do not fine-tune the bin edges any further. As we describe in Appendix C, the constraints for peaks (minima) are not very sensitive to the exact choice of the bin edges as long as at least a few bins are used.

D. Moments

We measure four low-order moments (two quadratic, one cubic, and one quartic). This set of moments is motivated by the perturbative expansion of MFs up to second-order in variance [30,63,65]. For a Gaussian random field, MFs can be expressed analytically using the variance of the field (σ_0) and the variance of the gradient of the field (σ_1). For non-Gaussian fields, one can write down an expansion of MFs in powers of the variance. The first-order term is described by the three skewness parameters and the second-order term by the four kurtosis parameters (see Refs. [63,65,66].) It has been shown that these perturbative series do not capture all the information that MFs are sensitive to for weak gravitational lensing fields unless the maps are considerably

¹²One could also study and compute the Minkowski tensors, a tensorial generalization of MFs [62] (see, e.g., <https://morphometry.org/>).

smoothed (e.g., smoothing scale $\sigma \geq 15$ arcmin) [30]. Since the tSZ field also traces late-time structure, we expect that these series do not converge for our maps using a Gaussian kernel with $\theta_{\text{FWHM}} = 1.4$ arcmin. We therefore treat the MFs and moments as separate descriptors in our study as done in previous works [31]. The un-normalized moments that we measure are the following:

$$\sigma_0 = \sqrt{\langle y^2 \rangle}, \quad \sigma_1 = \sqrt{\langle |\nabla y|^2 \rangle} \quad (12)$$

$$S_1 = \langle y^2 \nabla^2 y \rangle, \quad K_1 = \langle y^3 \nabla^2 y \rangle \quad (13)$$

As in the case of measuring the MFs, we only need the gradient and the Hessian of the tSZ field to compute the moments. Skewness terms vanish in the Gaussian case, while the kurtosis terms are nonzero. In this work, we measure the full kurtosis moment and do not subtract the Gaussian (“un-connected”) part. We note that the expansion of MFs up to the second-order includes two additional cubic terms and three additional quartic moments. We do not include them in our baseline results due to the large sample variance of those moments and the possible lack of convergence of their mean values with respect to the number of simulations that we use in this study. We explore the constraints from those additional moments in Appendix E.

IV. PARAMETER INFERENCE

To assess the constraining power of each of the summary statistics described in the previous section, we follow a standard Bayesian framework to infer the posterior distributions of the cosmological parameters $\theta_{i,j} = \{\Omega_c, \sigma_8\}$. We assume a Gaussian likelihood and a cosmology-independent covariance matrix [67,68] computed at the fiducial cosmology ($\theta_{\text{fid}} = \{0.264, 0.811\}$). Assuming flat priors and dropping the model-dependence, the posterior distribution is proportional to the likelihood, L

$$p(\theta|\mathbf{x}) \propto L(\theta) \propto \exp\left(-\frac{1}{2} \Delta \mathbf{x}^T \hat{C}_{\text{fid}}^{-1} \Delta \mathbf{x}\right). \quad (14)$$

Here, $\mathbf{x}$ is our observable (i.e., each tSZ statistic), $\Delta \mathbf{x}$ is the difference between the mean statistic predicted at a cosmology within the prior bounds and its fiducial value:

$$\Delta \mathbf{x} = \bar{\mathbf{x}} - \bar{\mathbf{x}}_{\text{fid}} \quad (15)$$

and $\hat{C}_{\text{fid}}^{-1}$ is the inverse covariance matrix:

$$C_{\text{fid}} = \frac{1}{N_{\text{fid}} - 1} \sum_{i=1}^{N_{\text{fid}}} (\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{x}_i - \bar{\mathbf{x}})^T. \quad (16)$$

The hat $\hat{}$ denotes that the covariance has been rescaled following Ref. [69] to remove bias in its estimation

$$\hat{C}_{\text{fid}}^{-1} = \frac{N_{\text{fid}} - d - 2}{N_{\text{fid}} - 1} C_{\text{fid}}^{-1}, \quad (17)$$

where d is the size of the observable vector $\Delta \mathbf{x}$ and $N_{\text{fid}} = 34,560$ is the number of simulations used to compute the covariance. In this work, d ranges between 1 and 300, depending on the exact set of statistics used to obtain the constraints and the binning choices. In all cases, the correction factor is very small (e.g., for the fiducial noiseless case when all of the descriptors are included $d = 157$ and the correction factor is ~ 0.995). We leave the exploration of using a cosmology-dependent covariance or a non-Gaussian form of the likelihood to future work.

Since our inference is in a 2D parameter space, we can easily compute the likelihood directly on a fine grid of the parameter values within the prior bounds. In this work, we use 1000×1000 linearly spaced points between $\Omega_c \in \{0.21, 0.32\}$ and $\sigma_8 \in \{0.6, 1\}$. At each point, we interpolate $\bar{\mathbf{x}}$ using a 2D Clough-Tocher cubic interpolator [70]. We construct the interpolator using the mean summary statistics, $\bar{\mathbf{x}}$, that have been computed over $N_{\text{sims}} = 5,112$ tSZ maps at each of the 153 cosmologies in the suite and $N_{\text{fid}} = 34,560$ at the fiducial. We can then determine the likelihood at each of the 10^6 points in the grid and obtain the confidence contours. We use the area of the 68% confidence contours, A_{68} , as the figure of merit in this work. We also compare the marginalized constraints on σ_8 . For this, we compute the mean parameters as

$$\theta_i^{\text{mean}} = \frac{\int L(\theta_i, \theta_j) \theta_i d\theta_i d\theta_j}{\int L(\theta_i, \theta_j) d\theta_i d\theta_j}, \quad (18)$$

and quote the errors by choosing the parameter values closest to the specified confidence limit of the marginalized likelihood.

V. RESULTS: NOISELESS MAPS

A. Cosmological dependence

In Fig. 5 we show the summary statistics computed for three example cosmologies over 5,112 maps: the fiducial and two others with Ω_c and σ_8 increased by 20% to illustrate their dependence on cosmological parameters. For the four moments, we list their values in Table II. For the moments, we compute their values over 34,560 maps for the fiducial cosmology in order to reduce their sample variance and determine accurate scalings with cosmological parameters.

First, we can clearly see the much stronger dependence of the observables on σ_8 compared to Ω_c .¹³ In the case of the

¹³For clarity, we therefore scale some of the ratios in Fig. 5 by a factor of 50.




FIG. 5. Summary statistics measured for three example cosmologies: the fiducial (green solid) cosmology and two cosmologies with Ω_c (blue dashed) and σ_8 (purple dot-dashed) increased by 20%. From left to right, the plots correspond to the power spectrum, three MFs (V_0 , V_1 , V_2), peaks, and minima. The stronger dependence on σ_8 is clearly evident. The bottom panels show the ratio between each of the two cosmologies and the fiducial. Variance on each of the observables is also shown.

power spectrum, the cosmological dependence is predominantly in the amplitude, which scales as $\propto \sigma_8^{7.4-7.7}$, in contrast to $\propto \Omega_c^{0.6-1.7}$ across the different multipoles. Similar dependence on the cosmological parameters has previously been shown, for example in Refs. [12,15,22]. Reference [22] found the power spectrum to scale as $\propto \sigma_8^{8.1} \Omega_m^{3.2}$.

For the four moments, we find scalings ranging from $\propto \sigma_8^{3.8-14.4}$ and $\propto \Omega_c^{(-0.2)-(+0.4)}$, where the range in the

TABLE II. Values and approximate cosmological parameter scalings for the four moments. The values are listed for the fiducial cosmology and two cosmologies with $1.2 \times \sigma_{8,\text{fid}}$ and $1.2 \times \Omega_{c,\text{fid}}$. The power-law scalings are defined as $\propto \sigma_8^{\alpha_{\sigma_8}}$ and $\propto \Omega_c^{\alpha_{\Omega_c}}$.

	$\sigma_0 \times 10^{-6}$	$\sigma_1 \times 10^{-7}$	$S_1 \times 10^{-19}$	$K_1 \times 10^{-23}$
Fiducial	1.54	1.27	-3.22	-1.94
$1.2 \times \sigma_8$	3.08	2.52	-24.2	-26.8
α_{σ_8}	3.8	3.8	11	14
$1.2 \times \Omega_c$	1.64	1.34	-3.35	-1.86
α_{Ω_c}	0.4	0.3	0.2	-0.2

powers reflects the different dependence for the four moments between σ_0 and K_1 as listed in Table II. As expected, the higher-order moments scale more steeply with parameters than the lower-order moments. Since σ_0 is the rms, the scalings for the variance, σ_0^2 , are similar to those found in Ref. [16] (i.e., $\sigma_0^2 \propto \sigma_8^{7.6}$ scales the same as the power spectrum, as expected). We note that K_1 scales negatively with Ω_c . A possible explanation for such dependence could be that in order to have a fixed σ_8 while increasing Ω_c (with h fixed), the amplitude of the primordial power spectrum, A_s , needs to decrease. The decrease in A_s leads to a reduction in the halo mass function that is larger than the amplification caused by the increase in Ω_c , which can lead to such negative scaling for some of the moments. For example, we have $A_s = 2.04 \times 10^{-9}$ at the fiducial cosmology with $\Omega_c = 0.264$ and $\sigma_8 = 0.811$, and $A_s = 1.64 \times 10^{-9}$ at the cosmology with $\Omega_c = 0.3168$ and $\sigma_8 = 0.811$. While we find that the mean values of the moments are converged within the standard error as a function of N_{sims} , the exact parameter scalings can vary due to some variation within the error bars. Therefore, we list the scalings computed using these example cosmologies and also show the scalings for each of the moments




FIG. 6. The left (right) panel shows the σ_8 (Ω_c) parameter scalings for the four moments computed using different sets of simulations: from Fisher simulations with parameters increased by 1% or 3% and from the example simulations, where parameters are increased by 20%. The errors are obtained by considering the lowest and highest scaling within the errors for a given moment.

computed using different sets of simulations in Fig. 6. The error bars here are obtained by taking the value of each moment at the fiducial and the example cosmology that leads to the smallest and largest scaling within the error of the given moment [e.g., in the case of positive moments, the lowest (highest) limit on the scaling can be obtained by taking the highest (lowest) fiducial value for the moment and the lowest (highest) value for the moment at a given cosmology].

We find that V_0 shows a noticeably different dependence on the two parameters, which leads to the degeneracy breaking that we discuss in subsequent sections. In particular, we see that Ω_c leads to a moderate, mostly constant increase in V_0 at high thresholds and a dip for a narrow range of low-significance thresholds. We find that σ_8 has the largest effect on the high-significance thresholds as well as a larger dip across a wider range of low-significance thresholds. This is similar to what has been found in Refs. [18,24] for the one-point PDF. V_1 and V_2 both exhibit a more pronounced cosmological dependence at the intermediate threshold values, as well as an overall steeper dependence at the high-significance thresholds. This suggests their higher sensitivity to Ω_c and σ_8 . As in the case of V_0 , the two cosmological parameters have different effects on V_1 and V_2 .

Finally, Ω_c and σ_8 affect different bins of the peak and minima histograms. We note that the peak histograms are generally more extended, covering a wider range of y values, while the minima histograms are narrower. The average number of peaks for the three cosmologies is $\approx 27,000$ – $29,000$, which is ≈ 240 – 260 peaks/deg². For the minima counts, the total number is of the same order: $\approx 26,000$ – $27,000$ or ≈ 240 minima/deg².

B. Constraints

In Fig. 7 (left), we show the credible contours obtained using noiseless tSZ maps with the power spectrum,

Minkowski functionals, peaks, minima, and moments. Note that we do not rescale the likelihood to a realistic f_{sky} in this section and only focus on the relative constraining power for each of the descriptors.¹⁴ We quantify the constraining power in terms of the 1σ contour area (A_{68}), the comparison to the power spectrum ($A_{68}^{\text{ps}}/A_{68}$), and the marginalized constraints on σ_8 ($\hat{\sigma}_8$). As can be seen in Fig. 7, the power spectrum does not have enough sensitivity to break the degeneracy between σ_8 and Ω_c , and thus the 1σ contour area for the power spectrum can be considered a lower bound.

We find that when all the observables are combined, the constraints are ~ 29 times tighter than from the power spectrum alone. The MFs yield tighter constraints than any of the other observables, while breaking the $\Omega_c - \sigma_8$ degeneracy. Compared to the power spectrum, the contour area is 21 times smaller. In fact, the constraints from all five statistics combined are only 1.4 times tighter than from the MFs alone.

Peak counts are the second-most constraining statistic. Peaks similarly break the $\Omega_c - \sigma_8$ degeneracy, but with overall wider credible contours in both the Ω_c and σ_8 directions, which yields 1.6 times tighter constraints than the power spectrum. Since the constraints from all the observables are tighter than from using the MFs alone, this suggests that peaks and MFs are complementary but still probe some of the same information.

Minima counts also break the degeneracy and have a similar tilt to the peak counts, but are less constraining in the direction of the best-constrained parameter combination. The contour area is actually larger than for the power spectrum ($A_{68}^{\text{ps}}/A_{68} = 0.84$). Since minima counts trace the “valleys” of the tSZ field, it makes intuitive sense

¹⁴For example, if we assume that the constraints would approximately scale by the square root of the sky area, the errors are expected to be ~ 12 times tighter for $f_{\text{sky}} = 0.4$.




FIG. 7. Left: 68% credible contours for the *noiseless* maps (smoothed with a Gaussian kernel with $\theta_{\text{FWHM}} = 1.4$ arcmin). The MFs yield $21\times$, and all statistics combined yield $29\times$, improvement over the power spectrum, as quantified via the area of the 1σ contours. Right: results obtained with each of the MFs separately. The improved constraints are driven by the combination of the three MFs, and dominated by V_1 and V_2 roughly equally.

why this statistic may contain less information in comparison to the peaks, for example.

Finally, we find that the four moments also achieve worse results compared to the power spectrum, $A_{68}^{\text{ps}}/A_{68} = 0.75$, and do not break the degeneracy between the two cosmological parameters. This is possibly due to the fact that the moments do not contain all the scale information captured by the power spectrum. The moments of gradients contain some spatial/scale information but likely not all of it. Additional scale information could further be included by combining moments measured for different smoothing scales, as we explore for the MFs in Sec. VI.

While the most striking effect of using MFs or combining all the summary statistics is the breaking of the $\Omega_c - \sigma_8$ degeneracy and a substantial shrinkage in the contour area, the marginalized errors on σ_8 are also considerably improved. With MFs, the error on σ_8 is almost 3 times smaller than the error obtained using the power spectrum. On the other hand, all other observables give σ_8 errors comparable to that from the power spectrum.

Looking at the constraints from each of the observables combined with the power spectrum, we see that using the power spectrum and MFs results in 23 times improvement over the power spectrum alone (as compared to the $21\times$ improvement found using the MFs alone, compared to the power spectrum). This suggests that the MFs do not contain all the information captured by the power spectrum and the two observables could be combined for better constraints. It is possible that the MFs are not sensitive to some of the

angular scales that we measure with the power spectrum. We find similar results for peaks, minima, and moments, which when combined with the power spectrum give ~ 3.4 , 1.9, and 1.2 times tighter constraints compared to using the power spectrum alone, respectively. This shows that each of these summary statistics is complementary to the power spectrum, probing different information in the tSZ maps. In particular, we note that the constraints from using the peaks and power spectrum are much tighter than using peaks alone. Similarly, combining minima counts and moments with the power spectrum achieves much better results than when each of those statistics is used individually. In general, this suggests that using the power spectrum in combination with one of these additional observables, in particular MFs or peak counts, achieves the best results.

In Fig. 8 (left), we show the correlation matrix for all the observables combined, defined as $\rho_{ij} = C_{ij}/\sqrt{C_{ii}C_{jj}}$, where C is the covariance. For clarity, we also show the correlation matrices for peaks, MFs, minima, and moments separately in the right panel. We can see the strong correlation across most multipoles for the power spectrum. Only the lowest bin has weak correlation with the rest. For the peaks, intermediate thresholds show the most correlation, while the second and the last bin contain independent information. Similar dependence is observed for the minima, except the correlations (and anticorrelations) are much stronger across the bins. V_0 shows strong correlations across the different bins. As expected, similar results were shown in Ref. [18] for the one-point PDF. As previously suggested, this is likely due to the same clusters




FIG. 8. Correlation matrix for the *noiseless* observables. The elements correspond to (1) power spectrum (77 bins), (2) peaks (9 bins), (3) MFs (60 bins), (4) minima (7 bins), and (5) moments (4 values). For the latter four observables, the correlation matrices are also shown separately on the right for clarity. All observables exhibit significant correlation across neighboring bins. The MFs show relatively low cross-correlation between the three MFs, and V_1 and V_2 have relatively low autocorrelation among their own bins. The cross-correlations between the different summary statistics are generally lower than the autocorrelations within each statistic.

contributing to multiple Compton- γ bins, as is most likely the case for the peaks and minima counts as well. V_1 and V_2 show less correlation between further-separated bins. The correlation is also low for some of the cross-elements between the three MFs. Finally, the moments show overall low correlation with one another. While there is significant correlation for the different bins for most of the observables, the cross-covariance between the different summary statistics, in particular between the power spectrum and the other descriptors, is mostly low. This is consistent with the results discussed above and the tight constraints that can be achieved by combining the various observables.

We summarize the 1σ contour areas for each of the statistics and the marginalized errors on σ_8 in Table III. We do not include the errors on Ω_c because the $\sigma_8 - \Omega_c$ degeneracy is not broken in all cases, including for the power spectrum, which serves as our baseline statistic.

C. Convergence with respect to the number of simulations

We check that we have a sufficient number of realizations for the covariance estimation by computing the contours using half of the simulations at the fiducial cosmology ($N_{\text{fid}} = 17,280$ for both the covariance and the mean summary statistic at the fiducial cosmology). With this number of simulations, the confidence area from all the summary statistics combined remains the same ($< 1\%$ difference). Note that some of the other summary statistics

require a larger number of simulations based on the individual covariance convergence tests described below (Fig. 9). However, this initial check shows that the constraints from all summary statistics combined, which give

TABLE III. Constraints obtained from different summary statistics on *noiseless* maps ($10.5 \times 10.5 \text{ deg}^2$ and smoothed with a Gaussian kernel $\theta_{\text{FWHM}} = 1.4 \text{ arcmin}$): 68% contour area (A_{68}), comparison to the power spectrum area ($A_{68}^{\text{ps}}/A_{68}$), and the mean value and marginalized error on σ_8 ($\hat{\sigma}_8$). We highlight the most significant improvement in the constraining power over the power spectrum in boldface.

Observable	$A_{68} \times 10^4$	$A_{68}^{\text{ps}}/A_{68}$	$\hat{\sigma}_8$
Power spectrum	7.82	1.0	$0.810^{+0.01}_{-0.009}$
Peaks	4.92	1.6	0.810 ± 0.012
MFs	0.367	21	$0.811^{+0.004}_{-0.003}$
Minima	9.31	0.84	0.811 ± 0.011
Moments	10.4	0.75	0.810 ± 0.011
All combined	0.267	29	0.811 ± 0.003
Power spectrum + peaks	2.28	3.4	0.810 ± 0.012
Power spectrum + MFs	0.346	23	$0.811^{+0.004}_{-0.003}$
Power spectrum + minima	4.11	1.9	0.811 ± 0.011
Power spectrum + moments	6.73	1.2	0.810 ± 0.011
V_0	6.21	1.3	$0.816^{+0.013}_{-0.019}$
V_1	2.16	3.6	0.810 ± 0.008
V_2	2.32	3.4	0.811 ± 0.006




FIG. 9. Convergence of the constraints with respect to the number of simulations. We show the ratio of the 68% confidence area obtained using a subset of simulations to the constraints obtained using all available simulations. Top: convergence of the mean summary statistic. Bottom: convergence of the covariance. Horizontal dashed lines mark $\pm 1\%$.

smooth contours, are stable to the number of simulations used at the fiducial cosmology to $\sim$ percent-level.

Next, we check that we have simulated enough realizations at each of the other cosmologies (i.e., that the mean statistics across the suite are converged). Although we have determined this via convergence tests performed prior to running the suite of simulations (see Appendix A), we can check this directly by computing the credible contour areas using a different number of simulations to compute the mean summary statistic at each cosmology (while keeping the number of simulations for the covariance estimation at $N_{\text{fid}} = 34,560$). For these tests, we compute the constraints using all of the summary statistics combined and each summary statistic individually. We plot the ratio between the 1σ area computed using all $N_{\text{sim}} = 5,112$ realizations at each cosmology and an area computed using some subset of the realizations in the upper panel of Fig. 9. We see that for all the observables, the contours are converged to within $\sim 1\%$ at 2,500 realizations and to within $\sim 2\%$ at 2,000 realizations.

Similarly, we can explicitly check the convergence as a function of the number of simulations used for the

covariance estimation (while using 5,112 maps for the mean, including the fiducial cosmology). We see that the covariance is converged for each of the summary statistics (lower panel of Fig. 9). For MFs, as few as $\approx 3,500$ are needed.

D. Minkowski functionals

To determine if the constraining power is driven by any specific MFs, we look at the credible contours for each of the MFs separately, which are shown in Fig. 7 (right). V_0 , V_1 , and V_2 all have distinct degeneracy directions. Thus, the tightest constraints are achieved when all three are combined. Conceptually, V_0 is equivalent to the one-point PDF, while V_1 and V_2 probe spatial information beyond what is captured in the “area.” It has been previously proposed that the one-point PDF is a near-optimal cosmological statistic for the tSZ field, since the tSZ field is dominated by the one-halo term and the one-point PDF captures information from all zero-lag moments up to arbitrarily high order [18,24]. However, here we see that the 1σ contour area is roughly $2.9\times$ and $2.7\times$ smaller for V_1 and V_2 , respectively, and $17\times$ smaller when all three MFs are combined, compared to using V_0 alone. One interpretation of these results is that each of the MFs are sensitive to a different range of halo masses, which leads to the different degeneracy directions for each of the MFs.

Next, we look at which MF bins contain the most sensitivity to cosmology. To do this, we compute credible intervals using only a subset of the thresholds over which we have measured the MFs (recall that we use 20 linearly-spaced bins between thresholds $\nu \in \{-1\sigma_{\text{fid}}, 3\sigma_{\text{fid}}\}$). In the case of V_0 , most of the constraining power lies in the lowest significance bins. In Fig. 10, we show the constraints computed for different subsets of the thresholds, from lowest to highest values, and show that the constraints are almost entirely coming from the lowest five bins (e.g., the area of the 1σ contour is just $\sim 20\%$ larger than when using all the thresholds). This can be explained by the large sample variance of the high threshold bins. Even though the tail of V_0 contains significant cosmological information as seen in Fig. 5, the sample variance reduces its constraining power. Note also how the degeneracy direction of the ellipse is different for each of the threshold bin subsets, suggesting that different parts of the halo mass function are probed.

In the case of V_1 and V_2 , the lowest half of the bins contains most of the sensitivity to cosmology. For these MFs, the full shape is important in capturing the information (e.g., the first 10 of the thresholds), while much less information lies in the high-significance tails. This is most likely due to the noticeably different degeneracy directions between the lowest bins and the rest (see Fig. 10). On the other hand, the intermediate and the high threshold bins seem to be probing similar information. From a practical




FIG. 10. Upper: constraints from each of the MFs using a subset of the thresholds over which each is measured (the fiducial threshold bins are between $\nu \in \{-1\sigma_{\text{fid}}, 3\sigma_{\text{fid}}\}$). For V_0 , almost all of the constraining power is in the five lowest significance thresholds, while for V_1 and V_2 , it is in the lowest half of the bins. Lower: each of the MFs at the fiducial cosmology. The vertical (dashed) lines and the shades of the markers correspond to the intervals considered for the constraints in the upper panel.

standpoint, this means that we do not necessarily need to measure the MFs at the highest thresholds. We can thus avoid possible numerical noise that can arise from the low values in those bins. Moreover, since the sample variance is the most difficult to capture at the variance levels corresponding to the most rare, massive clusters, possibly fewer simulations would be needed to obtain stable results with respect to the number of simulations and HMPDF halo mass function discretization settings when the threshold range is even more limited (e.g., to $\nu \in \{\nu_{\text{min}}, \nu_{\text{max}}\} \in \{-\sigma_{\text{fid}}, \sigma_{\text{fid}}\}$).

E. Comparison to Fisher forecast

In our analysis, we directly compute the posteriors for the cosmological parameters through the likelihood computation on a fine grid. This is possible via the suite of simulations that allow us to build an interpolator and predict the value of the summary statistics at any other cosmology. We compare our results to constraints obtained using the Fisher matrix formalism in order to:

- (1) Further cross-check that our results from the suite are converged.
- (2) Determine if Fisher forecasts may be sufficient for predicting the constraints from the MFs (given how close the constraints from the MFs are to an ellipse).

The Fisher matrix F_{ij} for the parameters $\theta_{i,j}$, and a likelihood L , can be written as follows

$$F_{ij} = \left\langle \frac{\partial \ln L}{\partial \theta_i} \frac{\partial \ln L}{\partial \theta_j} \right\rangle. \quad (19)$$

The marginalized errors can then be estimated via

$$\sigma_i = \sqrt{(F^{-1})_{ii}} \quad (20)$$

and the 68% Fisher ellipse area using

$$A_{\text{Fisher}} = \pi \Delta\chi^2 ab, \quad (21)$$

where a and b can be computed using Eq. (3) and $\Delta\chi^2 \approx 2.3$ for 68% CL. More precisely, to compute the Fisher matrix from simulations assuming a Gaussian likelihood, we need the derivatives of the mean summary statistic vectors ($\bar{\mathbf{x}}$) with respect to parameters $\theta_{i,j}$ and the inverse covariance matrix computed at the fiducial cosmology ($\hat{C}_{\text{fid}}^{-1}$) using $N_{\text{fid}} = 34,560$ simulations and rescaled with the Hartlap factor, as previously defined in Sec. IV. The Fisher matrix can then be calculated using

$$F_{ij} = \frac{\partial \bar{\mathbf{x}}}{\partial \theta_i} \hat{C}_{\text{fid}}^{-1} \frac{\partial \bar{\mathbf{x}}}{\partial \theta_j}. \quad (22)$$

To determine the derivatives, we consider three setups, where Ω_c and σ_8 are perturbed one at a time by 0.5%, 1%, and 3% from their fiducial values (a total of




FIG. 11. Comparison of the constraints from MFs obtained using the Fisher formalism versus using the full suite of simulations to compute the confidence contours as described in Sec. IV. From left to right, panels (1) and (2) show the convergence of the Fisher errors on Ω_c and σ_8 with respect to the number of simulations used to compute the derivatives: the ratio between the error computed using N_{sims} as indicated and all 5,112 simulations. The convergence is shown using Fisher cosmologies perturbed by 0.5%, 1%, and 3%. Horizontal black dashed lines mark $\pm 5\%$. Figure 3 shows the Fisher ellipse (computed for $\delta_\theta = 1\%$) and the credible contours computed using the full suite. The contour areas are in agreement to $\approx 2\%$.

12 cosmologies in addition to the fiducial). We calculate the derivatives from these simulations using a central difference scheme:

$$\frac{\partial \bar{\mathbf{x}}}{\partial \theta_i} \approx \frac{\bar{\mathbf{x}}|_{\theta=\theta_i+\delta\theta_i} - \bar{\mathbf{x}}|_{\theta=\theta_i-\delta\theta_i}}{2\delta\theta_i} \quad (23)$$

We compute the Fisher forecasts for the noiseless case, both for the MFs and the power spectra. It has been shown in previous works that simulation-based Fisher forecasts can give overly confident parameter error bars if the derivatives are not numerically converged, i.e., an insufficient number of simulations have been used [71,72]. To check for numerical convergence, we plot the 1σ errors as a function of the number of simulations in Fig. 11 for the MFs. We see good convergence for the 1% and 3% cases—the Fisher error bars are stable with respect to the number of simulations used to compute the derivatives. In the case of MFs, the Fisher constraints are also in very good agreement with the results obtained from using the full suite. In Fig. 11, we show the Fisher forecast ellipse using the 1%-variation cosmologies and the contours computed from the entire suite. The contour areas agree to $\approx 2\%$ from the two methods. While the marginalized errors are numerically converged for the 3%-variation simulations, the Fisher forecast using these simulations gives overly optimistic constraints (ellipse area is smaller by $\approx 15\%$ compared to the constraints computed using the full suite), suggesting the derivatives computed using these simulations are not accurate. As an additional numerical convergence check for the 1%-variation Fisher simulations, we split the simulations used to compute the derivatives into subsets of 1,700 and 3,412 simulations and compute the Fisher forecast using these two sets of simulations. From this comparison, we conclude that the contours computed from the full suite are consistent with Fisher forecasts. This suggests that in future work, we can use Fisher forecasts for the MF-based

forecasts instead of generating a full simulation suite if, for example, we want to compute the forecasts using more expensive simulations such as hydrodynamical simulations or N -body simulations with painted pressure profiles.

For the power spectrum case, we do not see the same numerical convergence with the number of simulations. One possibility is that the power spectrum has much less sensitivity to Ω_c , as is evident in the shape of the contours and the degeneracy in the $\Omega_c - \sigma_8$ direction. We experiment with discarding some of the power spectrum bins that lead to the noisiest derivatives but that does not improve the convergence of the Fisher errors. We compute the forecast for σ_8 alone and find good convergence both as a function of the number of simulations and across all $\delta\theta = 0.5, 1, 3\%$. The 1σ unmarginalized error on σ_8 is $\sigma_{\sigma_8} = 0.004$, consistent with the suite-based results.

F. Mass range

In addition to the statistical approach, the tSZ effect can be used to constrain cosmological parameters via cluster counts. It is therefore interesting to determine whether the tight constraints obtained via MFs are primarily driven by halos within a similar range of masses as those that would be detected in tSZ cluster surveys or by less massive halos, below the detection threshold. A careful comparison between cluster count cosmology analyses and the tSZ MFs, which would incorporate the subtleties regarding selection functions for cluster surveys, and the marginalization over unknown astrophysics (e.g., modeling the mass-observable relation for clusters or varying the pressure profile parameters in the tSZ map simulations) is outside the scope of this paper. Instead, we perform a simple test to see if the MF-based constraints are predominantly sourced by the most massive halos in our simulated maps.

To do this, we generate several sets of tSZ maps for which we restrict the maximum halo mass to a different




FIG. 12. Left: Fisher forecast constraints computed using simulations with a different maximum halo mass. The constraints tighten when the most massive halos are excluded (dark green). Lowering the maximum mass to $10^{13} M_\odot$ (cyan) weakens the 1D marginalized constraints, but does not affect the total Fisher ellipse area. Right: variance for each of the MF bins computed at the fiducial cosmology using simulations with the different mass cuts. The variance is reduced as the maximum mass is lowered, since the sample variance is highest for the most massive, rare halos.

value than the fiducial $M_{\max} = 10^{16} M_\odot$ and check how significantly the forecasts are impacted.¹⁵ Instead of simulating several additional full suites, we focus on a Fisher forecast comparison. As we have shown in the previous section, a Fisher forecast using MFs is in good agreement with our results from the full suite. Therefore, we choose to simulate tSZ maps for only five cosmologies: the fiducial model and four cosmologies with Ω_c, σ_8 perturbed by $\pm 1\%$ one at a time (see Sec. II for the exact Ω_c, σ_8 values).

We set the upper mass limit to $M_{\max} = 5 \times 10^{14} M_\odot$, motivated by the minimum mass in the latest ACT cluster catalog from Ref. [73]. As with the rest of the forecasts in this paper, we use 34,560 random realizations at the fiducial cosmology and 5,112 realizations at the perturbed cosmologies. In the upper panel of Fig. 12, we compare the Fisher constraints obtained using the full mass range in our fiducial analysis to the one with the reduced upper limit. We find that the constraints do not degrade by excluding the most massive halos. In fact, the Fisher errors somewhat tighten (the Fisher ellipse area is smaller by $\approx 18\%$). We check the convergence of the results by splitting the simulations used to compute the derivatives in the Fisher forecast into two separate batches (1,700 and 3,412 realizations) and computing the forecasts using each of the sets separately. We find that the constraints are in good agreement from both sets of maps, which means that we have a sufficient number of simulations to compute the derivatives. Similarly, we split the simulations at the fiducial cosmology in half and obtain the constraints using each separate half of the simulations for the covariance estimation. We find that our results are consistent.

¹⁵We choose to set a different $M_{\max}$ rather than computing the forecast for the most massive halos directly (i.e., setting $M_{\min}$ to a higher value) to avoid computing MFs over maps with a large number of zero-value pixels.

As additional tests, we compute Fisher errors using additional sets of maps with even more aggressive upper mass cuts of $M_{\max} = 10^{13} M_\odot$ and $10^{12} M_\odot$, which we also include in Fig. 12. In the $M_{\max} = 10^{13} M_\odot$ case, we see that the marginalized constraints for (Ω_c, σ_8) weaken by $\approx (58\%, 57\%)$, respectively. The total contour area, however, remains roughly the same ($< 1\%$ change). The Fisher ellipse stretches out in the Ω_c direction but tightens along the σ_8 direction. To cross-check this result, we compute the constraints adding a small amount of noise to the maps (white noise $0.01 \mu\text{K-arcmin}$ with no filters applied) as, realistically, the faint tSZ signal from low-mass halos will be washed out by noise. In this case, the overall constraints are indeed worse than those with $M_{\max} = 10^{16} M_\odot$. For $M_{\max} = 10^{12} M_\odot$, we lose roughly half of the constraining power and the errors on (Ω_c, σ_8) become larger by a factor of $\approx (3.2, 3.8)$. We note that the convergence of the Fisher errors requires more simulations for the cases with lower maximum halo mass cuts. We still confirm that the size of the errors converges to within $\approx 5\%$ as a function of the number of simulations used for computing the derivatives (i.e., the convergence test shown in Fig. 11).

We speculate that the improvement in the constraints when using $M_{\max} = 5 \times 10^{14} M_\odot$ is due to a decrease in the sample variance when removing the most massive objects in the maps. To show this, we plot the variances for each of the MF bins in the right panel of Fig. 12. It can be clearly seen that the MF variance decreases as we lower the maximum halo mass. As expected, we observe that the variance is very similar for the $M_{\max} = 10^{13} M_\odot$ and $10^{12} M_\odot$ cutoffs. We show the full absolute covariance matrices for the four maximum mass cases in Appendix D (see Fig. 21), which similarly show a reduction in the amplitude of the off-diagonal terms when using lower maximum mass cutoffs. We also include the correlation matrices, which show less correlation across the different MF bins as the halo mass upper limit is lowered.

TABLE IV. We compare the Fisher forecast constraints from MFs using simulations with different mass ranges. We list the maximum halo mass, A_F (the Fisher ellipse area), A_F^{16}/A_F (comparison to the fiducial constraints with $M_{\max} = 10^{16}M_\odot$), and the marginalized errors on Ω_c and σ_8 . We set $M_{\min} = 10^{11}M_\odot$ in all cases. For the noise in the fourth case, we use $0.01 \mu\text{K-arcmin}$ white noise and apply no filters to the maps. We find that in the idealized noiseless case, most of the constraining power of the MFs lies in the lower-mass halos, as the most massive objects contribute more to the variance than to the signal (see Sec. V).

$M_{\max} [M_\odot]$	$A_F \times 10^5$	A_F^{16}/A_F	σ_{σ_8}	σ_{Ω_c}
10^{16}	3.75	1.0	0.004	0.006
5×10^{14}	3.08	1.2	0.003	0.005
10^{13}	3.76	1.0	0.006	0.01
$10^{13} + \text{noise}$	5.63	0.7	0.007	0.011
10^{12}	7.42	0.5	0.014	0.02

These Fisher forecast tests with different maximum halo mass bounds have two main takeaways:

- (1) The constraints from MFs are driven primarily by the halo masses below the mass range that is individually detected in tSZ cluster surveys. In particular, the constraints improve by $\approx 1.2\times$ when we set the maximum mass to $M_{\max} = 5 \times 10^{14}M_\odot$, and remain roughly the same when we set $M_{\max} = 10^{13}M_\odot$. This suggests that the masses driving the MF constraints are lower than those that are directly detected for cluster cosmology analyses. Joint analyses using tSZ MFs and cluster counts, therefore, have the potential to achieve even tighter constraints on the cosmological parameters, since the two observables probe different halo mass ranges and can thus be highly complementary to each other. This conclusion is consistent with what has been learned about the tSZ power spectrum. For example, Ref. [20] has found that there is a lack of correlation between the high- ℓ tSZ power spectrum and cluster counts, as the two observables predominantly probe clusters at different redshift ranges.
- (2) Removing the most massive halos improves the constraints that can be achieved with the tSZ MFs by reducing the sample covariance, since the most massive objects have the highest sample variance. This again is consistent with the results found for the tSZ power spectrum (e.g., see Refs. [15,27,74–76]).

We summarize the results from these comparisons in Table IV. For each of the constraints, we measure the MFs using 20 bins between $\nu \in \{-\sigma_{\text{fid}}^x, 3\sigma_{\text{fid}}^x\}$ with the following variances: $\sigma_{\text{fid}}^{5 \times 10^{14}} = 1.1 \times 10^{-6}$, $\sigma_{\text{fid}}^{10^{13}} = 8.2 \times 10^{-8}$, $\sigma_{\text{fid}}^{10^{13} + \text{noise}} = 9.1 \times 10^{-8}$, $\sigma_{\text{fid}}^{10^{12}} = 1.1 \times 10^{-8}$.

VI. RESULTS: INCLUDING NOISE

A. Minkowski functionals

First, we study the impact of noise on the MFs, which are the most constraining summary statistic in the noiseless case.

1. Different experimental setups

We show the effects on the constraints of applying additional filters to the maps in panel (a) of Fig. 13. As described in Sec. II, we low- and high-pass filter the tSZ maps when combining with noise maps. We do this to avoid any problematic numerical behavior due to the high noise at the smallest scales and to avoid extrapolation beyond the ℓ range available in the public SO tSZ noise power spectra. Since we want the signal maps to contain the same angular scale range as the noise realizations, we apply the same filter to both the tSZ maps and the GRFs generated from the noise power spectra. As shown in panel (a) of Fig. 13, applying the filter increases the contour area by a factor of ≈ 2.3 due to the loss of modes. In the same plot, we also show the effects of applying a Fourier-space Wiener filter to noisy maps [see Eq. (5) in Sec. II]. To obtain the noisy constraints, we combine the smoothed tSZ maps with a GRF generated from the post-component-separation tSZ noise power spectra (more details are included in Sec. II). We apply this filter to up-weight the scales where the signal is more dominant and down-weight the scales where the noise is large. For the SO goal noise levels, the Wiener filter improves the constraints by a factor of ≈ 1.1 . For these contours, we compute the MFs using 30 bins between $\nu \in \{-3\sigma_{\text{fid}}^x, 3\sigma_{\text{fid}}^x\}$ with $\sigma_{\text{fid}}^{\text{SO goal}} = 1.9 \times 10^{-6}$ and $\sigma_{\text{fid}}^{\text{SO goal without } F_\ell^{\text{Wiener}}} = 3.2 \times 10^{-6}$. Adding noise makes the maps more Gaussian and thus the measured MFs are more symmetric around the zero-value threshold (i.e., the tails span similar numbers of negative and positive-value threshold bins) and broader compared to the noiseless cases in Fig. 5. We therefore choose to extend the range of thresholds over which we measure the MFs in order to capture their full shape.

Next, in panel (b) of Fig. 13, we show the results for three possible experimental setups: SO baseline, SO goal, and CMB-S4 noise levels. In each case, we have applied a corresponding Wiener filter. For these constraints, we again compute the MFs using 30 bins between $\nu \in \{-3\sigma_{\text{fid}}^x, 3\sigma_{\text{fid}}^x\}$ with the following variances: $\sigma_{\text{fid}}^{\text{SO,base}} = 2.1 \times 10^{-6}$, $\sigma_{\text{fid}}^{\text{SO,goal}} = 1.9 \times 10^{-6}$, $\sigma_{\text{fid}}^{\text{CMB-S4}} = 1.7 \times 10^{-6}$. With CMB-S4-level noise, the MF constraints are improved by $\approx 1.6\times$ compared to the constraints using SO goal noise. The SO baseline noise level gives constraints that are $\approx 0.9\times$ as tight compared to the SO goal noise level. These results are summarized in Table V.




FIG. 13. Constraints from MFs (V_0 , V_1 , and V_2) using various setups in comparison to the results for the noiseless maps (solid magenta in all panels). We emphasize that these results are computed for $10.5 \times 10.5 \text{ deg}^2$ maps, not maps covering the expected survey footprints of upcoming experiments. The panels show, starting from top left: (a) Effects of low- and high-pass filtering on the noiseless constraints and the impact of the Wiener filter on the noisy constraints for the SO goal experimental setup. (b) Constraints obtained using MFs on noisy maps for various experimental setups: SO baseline, SO goal, and CMB-S4. (c) Effects of combining MFs measured for different smoothing scales. (d) Constraints for various levels of white noise.

TABLE V. We compare the constraints from MFs for different experimental setups, as shown in panel (b) of Fig. 13. The results here are shown for $10.5 \times 10.5 \text{ deg}^2$ maps that have been low- and high-pass filtered, smoothed, combined with noise maps and Wiener-filtered as described in Sec. II. $A_{68}^{\text{SO goal}}$ corresponds to the 68% confidence area for the SO goal noise setup to which we compare the SO baseline and CMB-S4 constraints.

Noise level	$A_{68} \times 10^4$	$A_{68}^{\text{SO goal}}/A_{68}$	σ_8
SO baseline	9.18	0.9	0.810 ± 0.011
SO goal	7.92	1.0	0.810 ± 0.011
CMB-S4	5.07	1.6	0.810 ± 0.011

2. Smoothing scales

It has been shown in previous works that combining different smoothing scales can tighten the constraints from MFs (e.g., [29]). Smoothing has the effect of adding some spatial or scale information, especially to V_0 , which otherwise has none [29]. To test this on our noisy maps, we compute the MFs on two additional sets of noisy maps that have been additionally smoothed using a Gaussian kernel with $\theta_{\text{FWHM}} = 5 \text{ arcmin}$ and $\theta_{\text{FWHM}} = 20 \text{ arcmin}$. Note that in this case we smooth the total tSZ + noise map. For these results, we use the SO goal noise power spectra to generate the noise maps. We show the results in panel (c) of Fig. 13. The constraints are shown for:

TABLE VI. Constraints from MFs for the SO goal noise setup using different smoothing scales, as described in Sec. VI and shown in panel (c) of Fig. 13. The results here are shown for tSZ maps that are low- and high-pass filtered, smoothed with $\theta_{\text{FWHM}} = 1.4$ arcmin, combined with noise maps and either Wiener filtered or additionally smoothed using $\theta_{\text{FWHM}} = 5$ or 20 arcmin. The bottom four rows correspond to cases in which MFs measured from several setups are combined. We compare all the results to $A_{68}^{\text{SO goal}}$, which corresponds to constraints from noisy maps that only include the effects of the $\theta_{\text{FWHM}} = 1.4$ arcmin beam and a Wiener filter (row 1).

Smoothing scale(s)	$A_{68} \times 10^4$	$A_{68}^{\text{SO goal}}/A_{68}$	$\hat{\sigma}_8$
SO goal	7.92	1.0	0.810 ± 0.011
5'	22.6	0.35	$0.809^{+0.015}_{-0.014}$
20'	60.6	0.13	$0.804^{+0.036}_{-0.027}$
SO goal, 5'	7.82	1.01	0.810 ± 0.011
SO goal, 20'	7.83	1.01	0.810 ± 0.011
5', 20'	22.4	0.35	$0.810^{+0.015}_{-0.014}$
SO goal, 5', 20'	7.70	1.03	0.810 ± 0.011

- (i) tSZ + noise with a Wiener filter applied [denoted as SO goal, the same setup as the one in panels (a) and (b) in Fig. 13 in black solid curve];
- (ii) tSZ + noise with $\theta_{\text{FWHM}} = 5$ arcmin and $\theta_{\text{FWHM}} = 20$ arcmin smoothing applied to the composite map (denoted as 5' and 20');
- (iii) constraints from combining the different cases.

We again compute the MFs using 30 bins between $\nu \in \{-3\sigma_{\text{fid}}^x, 3\sigma_{\text{fid}}^x\}$ with $\sigma_{\text{fid}}^{\text{SO}, 5 \text{ arcmin}} = 2.0 \times 10^{-6}$ and $\sigma_{\text{fid}}^{\text{SO}, 20 \text{ arcmin}} = 0.9 \times 10^{-6}$. When we look at the two additional smoothing cases (5', 20') individually, simply smoothing the maps only weakens the constraints since we lose some of the small-scale information. The tightest constraints are from the fiducial case where simply a Wiener filter is applied. We find that combining the Wiener filter case with the 5 or 20 arcmin cases does not in any notable way improve the constraints over the Wiener filter constraints. In practice, combining scales increases the total dimension of the observable data vector and as such, the size of the covariance matrix and the number of simulations needed to properly estimate it. Our current results suggest that the gains from combining with 5 or 20 arcmin smoothing scales may not be large enough to be worth pursuing, but it may be worth testing smaller smoothing scales in future work. The 68% confidence areas for the different smoothing cases are listed in Table VI.

The tests using different smoothing scales indicate that much of the information captured by the MFs is lost if the maps are smoothed with a Gaussian kernel at $\theta_{\text{FWHM}} \gtrsim 5$ arcmin. In future works, it will be useful to explore additional smoothing scales, or different filters to determine which scales the MFs are most sensitive to.

3. White noise

Next, we compute the constraints for three different white noise levels and show the results in the last panel of Fig. 13. We perform this exercise to determine the approximate noise levels at which we recover close to noiseless results. We use flat (white) noise power spectra, as opposed to realistic tSZ post-component-separation noise, to be able to tune the exact noise amplitude in a simple way. For these tests, we generate the noise maps assuming a flat power spectrum (recall that the tSZ signal maps are already beam-convolved with a FWHM = 1.4 arcmin beam) and convert them from μK -arcmin to y-rad units at 150 GHz, as described in Sec. II. We low- and high-pass filter the maps and apply a Wiener filter. For the 0.1 and 1 μK -arcmin noise levels, we compute the MFs using 20 bins between $\nu \in \{-1.5\sigma_{\text{fid}}^x, 2.5\sigma_{\text{fid}}^x\}$, where $\sigma_{\text{fid}}^{0.1} = \sigma_{\text{fid}}^1 = 1.5 \times 10^{-6}$. For the 10 μK -arcmin noise level, we compute the MFs using 20 bins between $\nu \in \{-2\sigma_{\text{fid}}^{10}, 2\sigma_{\text{fid}}^{10}\}$ with $\sigma_{\text{fid}}^{10} = 1.2 \times 10^{-6}$. Again, we have chosen these ranges to capture the shape of the MFs since the maps are different for these three cases, as can be clearly seen in Fig. 20.

The credible contours in panel (d) of Fig. 13 show that the constraints are converging toward the results from the noiseless maps as the noise level decreases. This can also be visually seen in Fig. 20 and from the fact that the variance $\sigma_{\text{fid}}^{0.1} = \sigma_{\text{fid}}^1 = 1.5 \times 10^{-6}$ is similar to the noiseless case. The areas of the 68% contours are 0.883×10^{-4} , 1.14×10^{-4} , and 8.82×10^{-4} for the 0.1, 1, and 10 μK noise levels, respectively. The ratios between the 68% contour areas obtained using MFs computed for the noiseless maps and the different white noise levels are 0.4, 0.3, 0.04 (also listed in Table VII). For the 0.1 and 1 μK -arcmin noise levels, the contour areas are larger only by $\approx 3\%$ and $\approx 33\%$ compared to the noiseless case where a low- and high-pass filter has been applied [see panel (a) in Fig. 13]. This suggests that for upcoming experiments such as CMB-S4, which will have ≈ 1 μK -arcmin noise levels, improved component-separation techniques (e.g., via combination

TABLE VII. Constraints from MFs for various white noise levels, as described in Sec. VI and shown in panel (d) of Fig. 13. The results here are listed for $10.5 \times 10.5 \text{ deg}^2$ tSZ maps that have been smoothed and combined with white noise maps. We list the 68% contour area (A_{68}), comparison to noiseless constraints ($A_{68}^{\text{noiseless}}/A_{68}$), and the mean value and marginalized error on σ_8 ($\hat{\sigma}_8$).

Noise [μK -arcmin]	$A_{68} \times 10^4$	$A_{68}^{\text{noiseless}}/A_{68}$	$\hat{\sigma}_8$
0.1	0.883	0.4	$0.811^{+0.006}_{-0.005}$
1	1.14	0.3	0.811 ± 0.006
10	8.82	0.04	$0.810^{+0.012}_{-0.013}$




FIG. 14. 68% credible contours obtained using $10.5 \times 10.5 \text{ deg}^2$ noisy maps. Noise maps are generated using SO goal tSZ noise power spectra. The MFs give $1.6\times$ tighter constraints than the power spectrum. All summary statistics combined perform $1.8\times$ better than the power spectrum alone. The corresponding constraints are listed in Table VIII.

with CCAT-p data [77]) can push the constraints close to the noiseless scenario.

B. All summary statistics and realistic survey area

We show the noisy results for all the summary statistics in the left panel of Fig. 14. These are computed using the SO goal post-component-separation tSZ noise power spectra and include the Wiener filter (i.e., the setup matches the constraints shown for the MFs in the top two panels of Fig. 13, plotted in solid black). We find that in the noisy case, the MFs are still the most constraining observable ($\approx 1.6\times$ tighter constraints than when using the power spectrum alone). Combining the MFs and power spectrum yields $\approx 1.7\times$ improvement over the power spectrum alone. We find that the moments, peaks, and minima give worse constraints than the power spectrum. When combined with the power spectrum, peaks and minima result in $< 10\%$ improvement over the power spectrum, while moments can achieve $1.2\times$ tighter constraints than from the power spectrum alone. When we combine all the statistics, we achieve a $\approx 1.8\times$ improvement over the power spectrum alone. We note that the contours from all summary statistics combined are somewhat jagged and the relative improvement in constraining power can be $\approx 1.7\times$ if fewer cosmologies in the suite are used to compute the likelihood surface. Nonetheless, we see that the constraints are primarily driven by power spectrum + MFs with slight improvement from the other statistics. As the tSZ map becomes more Gaussian once the noise is added, the MFs contain less additional information. As described in

Sec. III, MFs can be written in terms of the moments for a weakly non-Gaussian field. Since we find that the MFs still perform better than the four moments and better than the combination of the full set of nine moments as described in Appendix E, it suggests that even in the presence SO-like levels of noise, the field is still highly non-Gaussian. This motivates a derivation of the analytic form of the MFs to higher-order corrections, which would allow determination of the exact set of moments needed to capture the information for the noisy tSZ maps, which could then be expressed analytically via, e.g., halo model tSZ polyspectra.

Examining each of the MFs individually, we find that V_1 and V_2 are notably more constraining than V_0 in the noisy case and that the combination of the MFs leads to the tightest results. In the noisy case, however, either using V_1 or V_2 results in constraints close to what can be achieved from combining all three ($\approx 1.4\times$ or $1.5\times$ versus $\approx 1.6\times$ improvement over the power spectrum). We find that peaks, minima, and moments are not as constraining in the noisy case. This could possibly be due to noise washing out the peaks that are most sensitive to cosmology and Gaussianizing the field. Still, we find that combining the power spectrum with moments improves the constraints, suggesting that the two are complementary.

We show the correlation matrix for the noisy maps in Fig. 15. As expected, the noise significantly reduces the correlation across the different bins for most of the observables. We summarize the constraints for all the summary statistics in Table VIII.

Finally, we scale the results from the power spectrum, MFs, and moments to a realistic sky fraction $f_{\text{sky}} = 0.4$ to approximately quantify the size of the parameter error bars that could be obtained from surveys like SO. Note that we do not include foregrounds in our simulated maps, apart from the residual foreground power in the post-component-separation tSZ noise power spectra, or marginalize over pressure profile parameters in this analysis. Therefore, these scaled forecasts represent an approximate idealized upper limit on how well we could constrain cosmological parameters with tSZ statistics.

The noisy posteriors hit the prior bounds, as shown in Fig. 14. This makes scaling the marginalized errors on Ω_c and σ_8 computed directly from these posteriors not entirely accurate. Thus, instead, we compute the marginalized error on the best-constrained parameter combination $\Sigma_8 = \sigma_8(\Omega_c/0.264)^\alpha$, where α is the slope of the line passing through the constraints, determined from the parameter covariance. The degeneracy directions vary for the observables: we find $\alpha = 0.07$ and 0.11 for the power spectrum and MFs, respectively. Therefore, it is best to compare how constraining each of the summary statistics is along its best-constrained parameter combination. We find $\Sigma_8^{\text{ps}} = 0.810^{+0.006}_{-0.005}$, and




FIG. 15. Correlation matrix for the tSZ statistics measured on *noisy* maps (the associated parameter constraints are shown in Fig. 14). The elements correspond to (1) power spectrum (77 bins), (2) peaks (9 bins), (3) MFs (90 bins), (4) minima (7 bins), and (5) moments (four values). Noise reduces the correlation between the bins for the power spectrum, peaks, MFs, and minima, as well as the cross-covariance between the different observables.

$\Sigma_8^{\text{MFs}} = 0.810 \pm 0.004$. Taking the average of the error bars (due to asymmetry), the MFs give error bars tighter by $\approx 33\%$ than the power spectrum. Dividing the errors on Σ_8 by $\approx \sqrt{16501 \text{ deg}^2 / 110 \text{ deg}^2} \approx 12.2$, we find that the

TABLE VIII. Constraints obtained from different summary statistics on *noisy* maps ($10.5 \times 10.5 \text{ deg}^2$): 68% contour area (A_{68}) in the $\Omega_c - \sigma_8$ plane, comparison to the power spectrum contour area ($A_{68}^{\text{ps}}/A_{68}$), and the mean value and marginalized error on σ_8 ($\hat{\sigma}_8$). For the noise maps, we use the SO goal tSZ noise power spectra, and apply both a low-/high-pass filter and a Wiener filter.

Observable	$A_{68} \times 10^4$	$A_{68}^{\text{ps}}/A_{68}$	$\hat{\sigma}_8$
Power spectrum	12.3	1.0	$0.810^{+0.01}_{-0.009}$
Peaks	38.6	0.32	0.809 ± 0.023
MFs	7.92	1.6	0.810 ± 0.011
Minima	31.0	0.40	0.809 ± 0.018
Moments	12.2	1.01	$0.810^{+0.011}_{-0.01}$
All combined	6.78	1.8	$0.810^{+0.01}_{-0.009}$
Power spectrum + peaks	11.7	1.1	0.810 ± 0.023
Power spectrum + MFs	7.23	1.7	0.810 ± 0.011
Power spectrum + minima	11.9	1.04	0.810 ± 0.018
Power spectrum + moments	10.6	1.2	$0.810^{+0.01}_{-0.009}$
V_0	28.4	0.43	0.810 ± 0.019
V_1	8.48	1.5	$0.810^{+0.012}_{-0.013}$
V_2	9.02	1.4	0.811 ± 0.012

MFs can constrain Σ_8 to $\approx 0.04\%$ precision, while the power spectrum alone can place an $\approx 0.06\%$ constraint. We emphasize again that the proper inclusion of foreground contamination and marginalization over ICM astrophysics parameters will weaken these constraints, but nevertheless it is clear that the tSZ field contains significant cosmological information.

VII. CONCLUSIONS

In this work, we have studied the cosmological information content of several tSZ summary statistics: the power spectrum, MFs, peaks, minima, and moments. To do this, we have generated a large suite of simplified tSZ simulations (more than 800,000 maps of size $10.5 \times 10.5 \text{ deg}^2$ each, with 0.1 arcmin pixels) with varying cosmological parameters Ω_c and σ_8 . This is the first systematic study of several beyond-Gaussian tSZ statistics to our knowledge. Our main findings can be summarized as follows:

Noiseless maps

- (1) MFs significantly outperform all the other summary statistics. Using MFs, we can achieve $\sim 21\times$ tighter constraints compared to those obtained from the tSZ power spectrum, as quantified by the area of the 68% contour in the $\Omega_c - \sigma_8$ plane. The constraints are driven by the combination of the three MFs—mostly by V_1 and V_2 —due to their different degeneracies. Thus, we find that the MFs contain substantial additional information beyond the one-point PDF (i.e., since V_0 contains the information from the

one-point function of the field). Combining all five summary statistics leads to a ~ 29 -fold improvement over the constraints from the power spectrum alone. Combining each of the four summary statistics with the power spectrum yields $\approx 23, 3.4, 1.9, 1.2\times$ improvement for MFs, peaks, minima, and moments, respectively, over using the power spectrum alone.

- (2) We compute a Fisher forecast for several additional sets of simulations, which have a lower maximum halo mass cutoff: $M_{\text{max}} = 5 \times 10^{14}, 10^{13}, 10^{12} M_{\odot}$. We find that the MF-based constraints are improved by $\approx 1.2\times$ when we exclude the most massive halos, i.e., those that would be individually detected in tSZ cluster surveys ($M_{\text{max}} = 5 \times 10^{14} M_{\odot}$), while the marginalized errors on Ω_c and σ_8 tighten by $\approx 12\%$ and 8% , respectively. Such improvement in the constraints is likely due to a reduction in the sample variance from excluding the most massive halos. This suggests that combined analyses using tSZ MFs and cluster counts may provide powerful constraints.

Noisy maps

- (1) In the noisy case, MFs again yield the best constraints, outperforming the tSZ power spectrum by a factor of ~ 1.6 . Combining the five summary statistics gives a factor of ~ 1.8 improvement over the power spectrum alone. Combining each of the four summary statistics with the power spectrum gives $\approx 1.7, 1.2, 1.1, 1.04\times$ tighter constraints for MFs, moments, peaks, and minima, respectively, over using the power spectrum alone. We note that recently Ref. [40] has similarly shown that MFs achieve the tightest constraints compared to peaks, minima, and the one-point PDF, when applied to weak gravitational lensing convergence maps from the Subaru Hyper Suprime-Cam first-year data.
- (2) For the MFs, we show that the results approach the noiseless case for white noise levels of 0.1 or $1 \mu\text{K-arcmin}$ at 150 GHz . We also demonstrate that combining MFs measured with different smoothing scales does not improve the constraints.
- (3) For a realistic $f_{\text{sky}} = 0.4$, we find that the MFs and power spectrum respectively achieve an $\approx 0.04\%$, and 0.06% error bar on the best-constrained parameter combination $\Sigma_8 = \sigma_8(\Omega_c/0.264)^\alpha$. Here, $\alpha = 0.11$ and 0.07 for the MFs and power spectrum, respectively.

We note that there are several simplifying assumptions in this analysis that would be important to study in follow-up work:

- (1) *Varying astrophysical parameters:* We fix all parameters associated with the ICM gas pressure profiles and focus only on the sensitivity to cosmological parameters. In practice, we would need to jointly fit

the astrophysical parameters and cosmological parameters, since the ICM pressure profile is currently not well-constrained over the full mass and redshift range of interest here. It would, therefore, be interesting and important to see how well these summary statistics can perform when astrophysical parameters are allowed to vary and if they can constrain some of the astrophysical parameters just as effectively. For example, Ref. [16] used the different dependence of the variance and skewness on cosmological and astrophysical parameters to construct a statistic that cancels out the dependence on cosmology and can thus be used to constrain the astrophysical parameters. We can test this breaking of the degeneracies between ICM physics and cosmology by applying some of the statistics studied in this work.

- (2) *Effects of clustering:* We use simplified tSZ simulations that do not account for realistic clustering of halos in this work. Although clustering effects are expected to be small for tSZ statistics [15,24,45], it would be important to quantify this via a direct comparison with hydrodynamical or N -body simulations.
- (3) *Realistic noise:* Another simplifying assumption in our noisy results is that we model both the foreground residuals in the tSZ maps after component separation and instrumental noise as GRFs using the estimated tSZ noise power spectra. In practice, there will be residuals in the maps that are non-Gaussian and correlated with tSZ signal (e.g., cosmic infrared background and point sources). In order to properly account for these effects, it will be important to forward model the astrophysical foregrounds and tSZ signal, processing the simulated maps through a component separation pipeline.
- (4) *Gaussian likelihood assumption and constant covariance:* In this work, we assume a Gaussian likelihood for all observables. We have performed some checks on this assumption. As an example, in the noiseless case, the distribution of the signal realizations at the fiducial cosmology in most MF bins are close to Gaussian-distributed. Similarly, Refs. [78–80] have shown that the Gaussian approximation is sufficient for many WL NG statistics. Nevertheless, it is important to directly quantify the effects of the Gaussian likelihood assumption on the tSZ constraints via, for example, simulation-based inference frameworks (e.g., [40,81]). Since the tSZ simulations in this work are very inexpensive, such a comparison is not impeded by the availability of simulations and can thus be studied in great detail. Similarly, the effects of using a cosmology-dependent covariance should be quantified in future work.

In summary, our results show that there is significant non-Gaussian information in the tSZ field, both in the noiseless and noisy cases. We find that the majority of this constraining power comes from halos that are not likely to be individually detected in upcoming tSZ cluster surveys, so the derived cosmological constraints are complementary to those found from tSZ cluster counts. This work is an important step toward characterizing the full information content of the tSZ field, with many interesting avenues for follow-up work.

In addition to the software listed in the paper, we also make use of NumPy [82], SciPy [70], Matplotlib [83], PIXELL,¹⁶ and COLORGRICAL [84] throughout this work.

We make some data used in this analysis available at [85]:

- (1) The simulations at the fiducial cosmology and those with Ω_c, σ_8 perturbed by $\pm 1\%$ for Fisher forecasts. We share the full $f_{\text{sky}} = 0.1$ maps at the simulated 0.1 arcmin resolution (960 maps for the fiducial model and 142-144 for the four others).
- (2) Fiducial noiseless and noisy summary statistic data vectors ($816, 696 \times 5 \times 2 = 8, 166, 960$ files).

We also make the main analysis pipeline available at [86]. Due to the large storage requirements for the full simulation suite, all other simulations can be provided upon reasonable request to the corresponding author.

ACKNOWLEDGMENTS

We are very grateful to Leander Thiele for assistance with HMPDF and Boris Bolliet for help with CLASS_SZ. We would also like to thank Shivam Pandey, Greg Bryan, Lam Hui, Sam Goldstein, Boris Bolliet, and Oliver Philcox for useful suggestions and discussions. We thank the anonymous referee for useful comments. A. S. and J. C. H. acknowledge support from NASA Grant No. 80NSSC23K0463 (ADAP). Z. H. and J. C. H. acknowledge support from NASA Grant No. 80NSSC24K1093 (ATP). J. C. H. also acknowledges support from the Sloan Foundation and Simons Foundation. We acknowledge the use of the NSF XSEDE facility for the simulations and data analysis in this study. We acknowledge 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. This is not an official Simons Observatory Collaboration paper.

DATA AVAILABILITY

The data that support the findings of this article are openly available [85,86].

APPENDIX A: NUMERICAL CONVERGENCE OF TSZ MAPS

Our tSZ simulations draw the number of halos from a Poisson distribution by discretizing the halo mass function into bins of redshift and mass. Since the tSZ signal is dominated by the most massive halos, we must determine how many simulations are needed to capture the associated sample variance and if the mass function discretization settings are sufficient, before running the suite of different cosmologies. To do this, we simulate 5,184 maps using different choices for the number of bins in redshift (N_z) and mass (N_M) and look at the mean of each of the five summary statistics with respect to the number of simulations, N_{sims} .

The top panel in Fig. 16 shows the angular power spectrum values for 3 of the 159 linearly spaced ℓ bins centered between $\ell = \{50, 7950\}$, which are consistent between $N_z = N_M = 150$ and $N_z = N_M = 200$ at $N_{\text{sims}} \sim 5,000$ within the standard errors.

In the second and third panels, we show the results for peak and minima counts. We determine the fiducial peak-count bin definitions by taking a realization of a single tSZ map and binning the peaks into 10 bins in such a way that every bin has an equal number of peaks. We merge the two highest-significance bins to ensure the peak counts are consistent within the standard errors for the three highest-precision numerical settings and end up with a total of 9 bins in our fiducial setup. We check the three highest-precision numerical settings for this bin because of the more significant variation across different numerical settings compared to other observables. We show the first, fifth, and the last histogram bin in the figure. We see that the peak counts for the highest bin are significantly different when low values for N_z and N_M are used, but converge to the same mean for the higher values of N_z and N_M .

Similarly, the bottom panel includes the results for minima counts. Since minima represent “valleys” in the tSZ field rather than possibly individual halos, we do not see such substantial effects. We follow a similar approach of determining 10 bins with equal number of minima and merging the bins to ensure convergence of the mean between $N_z = N_M = 150$ and $N_z = N_M = 200$. We adopt 7 bins in our fiducial setup, three of which are shown in the figure. To further illustrate the effects of numerical discretization, in Fig. 17 we plot the peak and minima counts for 100 linearly spaced bins between $\{-0.75, 3\}\sigma_{\text{fid}}$ and $\{-1, 0.5\}\sigma_{\text{fid}}$, respectively. The peak histogram clearly shows wiggles in the high-significance tail due to insufficient discretization at low values of N_z and N_M .

¹⁶<https://github.com/simonsobs/pixell>




FIG. 16. Convergence of the mean of the different statistics with respect to the number of simulations and discretization of the halo mass function (i.e., N_z and N_M correspond to the number of bins in redshift and mass, respectively). Top: angular power spectrum at $\ell = 200, 1000, 7000$ as a function of the numerical settings. Middle: the convergence of some of the fiducial bins for peak counts. Note the discrepancy in peak counts for low N_z and N_M . Bottom: convergence of the minima counts.




FIG. 17. We show the finely binned peak (minima) histograms in the left (right) panel (100 linearly spaced bins). Note the numerical wiggles in the high-significance peak count bins when the halo mass function is not sufficiently discretized (i.e., when using low N_z and N_M).




FIG. 18. Convergence of the mean value of MFs and moments with respect to the number of simulations and the discretization of the halo mass function (i.e., N_z and N_M correspond to the number of bins in redshift and mass, respectively).




FIG. 19. Validation of the simplified tSZ simulations via power spectrum comparisons. The black (gray) dashed lines indicate 1% (2.5%) difference levels. Left: The ratio between the average of 300 power spectra computed from maps with different resolutions to the power spectrum from the maps at 0.05 arcmin resolution, when smoothed using a Gaussian kernel with $\theta_{\text{FWHM}} = 1.4$ arcmin. Right: Comparison of the estimated power spectra for our fiducial cosmology to that computed from the halo model code CLASS_SZ. The power spectra are estimated using 5,112 realizations of $10.5 \times 10.5 \text{ deg}^2$ maps.

For the MFs (Fig. 18), we show one bin for each of V_0 , V_1 , and V_2 . We check that all the bins for each V_0 , V_1 , and V_2 are consistent between $N_z = N_M = 150$ and $N_z = N_M = 200$ within one or two standard errors. We also show the convergence for some of the moments (Fig. 18) and check that they are converged to the same mean. Based on these results, we adopt the numerical settings $N_z = N_M = 150$ and conclude that ~ 5000 maps suffices for each cosmology.

Additionally, we compare the power spectra computed from the maps generated at different resolutions to

determine the lowest map resolution that we can use without losing the accuracy of the signal, given the $\theta_{\text{FWHM}} = 1.4$ arcmin beam smoothing, which we apply in all the forecasts. The left panel in Fig. 19 shows the ratios for power spectra computed from several sets of maps at different resolutions compared to the power spectrum at 0.05 arcmin resolution. For this comparison, we use 300 realizations of $\sim 10 \times 10 \text{ deg}^2$ maps, remove the square pixel window function, and multiply each power spectrum by the mean of the apodization mask, or $\sum_{n=1}^N M(n)/N$,




FIG. 20. Simulated tSZ map combined with a Gaussian random field generated from a flat (white) power spectrum with noise levels corresponding to 10, 1, and 0.1 $\mu\text{K-arcmin}$ at $\nu = 150$ GHz. In the cases of 1 or 0.1 $\mu\text{K-arcmin}$ white noise, the field closely resembles the noiseless map.

where M is the apodization mask value at each pixel n and N is the total number of pixels in the given map (i.e., approximating that there is an equal suppression of power across all scales due to the apodization mask). We have checked that this simple multiplicative factor is sufficient in this case by comparing to the power spectra estimated with the full MASTER pseudo- C_ℓ approach for masked fields implemented in NaMaster [87].¹⁷ Based on the results in Fig. 19, we conclude that 0.1 arcmin resolution is sufficient.

To check the accuracy of the simulated tSZ signal, we compare the power spectrum from the maps against the halo-model code CLASS_SZ [12,22,27,52,53]. In this comparison, we show the results averaged over 5,112 map realizations. We divide out the pixel window function, account for the apodization mask, and multiply the theoretical halo model prediction by the same Gaussian beam. The power spectra are in good agreement (close to percent level at most ℓ) as shown in the right panel of Fig. 19.

APPENDIX B: WHITE NOISE

In Fig. 20, we show tSZ maps combined with different levels of white noise (see Sec. II for a description).

APPENDIX C: ADDITIONAL NOISELESS CONSTRAINTS TESTS

In addition to testing the stability of the noiseless constraints to the number of simulations used for the covariance or mean summary statistic estimations described in Sec. VC, we test the constraints against interpolation errors and fiducial binning choices.

We check that the errors from interpolating the summary statistics at the cosmologies where we do not have simulations are negligible. We do this by testing the accuracy of the interpolator by excluding the fiducial

cosmology in its construction (i.e., using 153 cosmologies) and checking the difference in the predicted observable values versus those measured directly from the fiducial cosmology. We find that the difference between interpolated and true values are all $< 0.03\sigma$ with a median of 0.01σ , where σ refers to the rms in each observable bin. We also check that the contours remain the same if we only use a subset of cosmologies. For example, we compute the contours using only 102 cosmologies from the initial LH suite (see Sec. II) and confirm that the contours remain the same (i.e., $\sim 1\%$ difference in contour area in the case of using all summary statistics combined). With these tests, we conclude that any errors from interpolation are minor.

We check how stable our results are to binning choices for the MFs, peaks, and minima. For peaks and minima, we do so by computing the constraints using different bin edges. For this exercise, we keep the outermost bins the same (e.g., $[-100, -0.3, \dots, 1, 1000]\sigma_{\text{fid}}$), while combining or splitting the rest of the bins. We find that for peaks, the results remain stable and almost the same regardless of the number of bins ($< 1\%$ difference using 30 or 16 versus 9 bins, and $\sim 7\%$ change when using 4 bins). The contour area is only significantly impacted when we use a single bin for the peak counts. For the minima counts, using 4, 12, or 22 bins similarly leads to small differences in the area: $\sim 18\%$, 3% , and 4% increase, respectively. The contour sizes are most stable for 12 and 22 bins. For the MFs, we find that the areas are stable when increasing the number of bins to 50, 80, or 100. From these tests, we conclude that the constraints presented in Sec. V are stable to these choices.

APPENDIX D: DIFFERENT MAXIMUM MASS CUTOFFS: COVARIANCE AND CORRELATION MATRICES

In Fig. 21, we show the correlation and absolute values of the covariance matrices for the MFs computed from simulations with different maximum halo mass cutoffs, as

¹⁷<https://github.com/LSSTDESC/NaMaster>




FIG. 21. Top: correlation matrices computed for the fiducial cosmology (34,560 realizations) using simulations with different maximum halo mass cutoffs. Excluding massive halos reduces the correlation between the MF bins. Bottom: absolute covariance matrices for the different mass cutoffs. The covariance is similarly reduced as the maximum halo mass used in the simulation is lowered.

studied in Sec. V and in Fig. 21. Both correlation across the MF bins and the total covariance are lowered when the most massive halos are excluded from the maps.

APPENDIX E: ADDITIONAL MOMENTS

The full set of moments required to write down the perturbative expansion of the MFs up to second order in variance consists of nine moments. In the main text, we have restricted our results to a total of four moments by including only one cubic and one quartic moment due to possible lack of convergence of the rest of the cubic and quartic moments (S_0 , S_2 , K_0 , K_2 , K_3):

$$S_0 = \langle y^3 \rangle, \quad S_2 = \langle |\nabla y|^2 \nabla^2 y \rangle \quad (\text{E1})$$

$$K_0 = \langle y^4 \rangle, \quad K_2 = \langle y \nabla^2 y |\nabla y|^2 \rangle, \quad K_3 = \langle |\nabla y|^4 \rangle \quad (\text{E2})$$

In Fig. 23, we show an example convergence plot for one of these moments, K_2 , for three cosmologies: fiducial and those with σ_8 and Ω_c increased by a factor of 1.2. In these plots, we show the convergence of the mean value using 34,560 simulations for the fiducial and 15,552 for the example cosmologies. The plot shows that K_2 value is likely unconverged using 5,112 sims at each of the cosmologies in the suite. For completeness, we include a table with the values for the rest of these moments at the




FIG. 22. Same as Fig. 6 for the additional two cubic and three quartic moments.

TABLE IX. Same as Table II, but for the five additional moments not considered in the main text.

	$S_0 \times 10^{-17}$	$S_2 \times 10^{-22}$	$K_0 \times 10^{-21}$	$K_2 \times 10^{-26}$	$K_3 \times 10^{-25}$
Fiducial	5.74	-8.19	4.18	-7.52	1.10
$1.2 \times \sigma_8$	41.7	-60.9	53.1	-97.5	13.9
α_{σ_8}	10.9	11.0	13.9	14.1	13.9
$1.2 \times \Omega_c$	6.07	-8.59	4.06	-7.69	1.15
α_{Ω_c}	0.3	0.3	-0.2	0.1	0.2




FIG. 23. Convergence of the K_2 moment as a function of the number of simulations for three different cosmologies: fiducial and two example cosmologies with σ_8 and Ω_c parameters increased by $1.2\times$. We show the mean at each number of simulations (dashed gray), and the mean at the maximum number of simulations, $N_{\text{fid}} = 34,560$ and $N_{\text{sim}} = 15,552$ (black solid). The vertical line corresponds to 5,112 simulations, which is the number of simulations used for the mean value of the summary statistics at all cosmologies except the fiducial.

TABLE X. Summary of the constraints obtained using nine moments in both the noiseless and noisy cases.

Observable	$A_{68} \times 10^4$	$A_{68}^{\text{ps}}/A_{68}$	$\hat{\sigma}_8$
Noiseless			
σ_0, σ_1	16.9	0.46	$0.81^{+0.011}_{-0.01}$
σ_0, σ_1, S_0	16.9	0.46	$0.81^{+0.011}_{-0.01}$
$\sigma_0, \sigma_1, S_0, S_1$	12.6	0.62	$0.81^{+0.011}_{-0.01}$
$\sigma_0, \sigma_1, S_0, S_1, S_2$	12.5	0.63	0.81 ± 0.01
$\sigma_0, \sigma_1, S_0, S_1, S_2, K_0$	12.1	0.65	0.81 ± 0.01
$\sigma_0, \sigma_1, S_0, S_1, S_2, K_0, K_1$	9.96	0.79	$0.81^{+0.011}_{-0.01}$
$\sigma_0, \sigma_1, S_0, S_1, S_2, K_0, K_1, K_2$	9.94	0.79	$0.81^{+0.011}_{-0.01}$
$\sigma_0, \sigma_1, S_0, S_1, S_2, K_0, K_1, K_2, K_3$	9.86	0.79	0.81 ± 0.01
Noisy			
σ_0, σ_1	18.3	0.67	$0.809^{+0.012}_{-0.01}$
σ_0, σ_1, S_0	18.3	0.67	$0.809^{+0.012}_{-0.01}$
$\sigma_0, \sigma_1, S_0, S_1$	14.1	0.87	$0.810^{+0.011}_{-0.01}$
$\sigma_0, \sigma_1, S_0, S_1, S_2$	13.2	0.93	0.810 ± 0.011
$\sigma_0, \sigma_1, S_0, S_1, S_2, K_0$	13.2	0.93	$0.810^{+0.011}_{-0.01}$
$\sigma_0, \sigma_1, S_0, S_1, S_2, K_0, K_1$	11.9	1.04	$0.810^{+0.011}_{-0.01}$
$\sigma_0, \sigma_1, S_0, S_1, S_2, K_0, K_1, K_2$	11.2	1.1	0.810 ± 0.011
$\sigma_0, \sigma_1, S_0, S_1, S_2, K_0, K_1, K_2, K_3$	11.1	1.1	0.810 ± 0.011

three cosmologies and a plot of their scalings with cosmological parameters in Table IX and Fig. 22, similar to what we have shown in Sec. V.

In Table X, we summarize the results we obtain when we include these additional moments both in the baseline noiseless and noisy cases. We also show the constraints for each moment individually, except σ_0 , which is not constraining on its own. In the noiseless case, we find that by including all nine moments, we still obtain constraints weaker than those from the power spectrum, with $A_{68}^{\text{ps}}/A_{68} = 0.79$. But the constraints are $\approx 1.1\times$ tighter than when using only the four moments considered in the main text. In the noisy case, we find that the constraints from all moments are $\approx 1.1\times$ better than the power spectrum alone and $\approx 1.1\times$ better than when using just the four moments.

Looking at the moments individually, we see that in the noiseless case, the constraints are improved the most by inclusion of each new higher-order moment of the derivatives (e.g., when S_1 or K_1 are added), while adding additional higher-order moments with or without the derivatives has a less substantial effect. In the noisy case, we see similar behavior, but adding additional higher-order moments of the derivatives has larger effects (e.g., adding S_2 or K_2). In the future, it may be advantageous to explore various ways of Gaussianizing the tSZ field by applying a nonlinear transformation to the maps before measuring these moments to more easily reach convergence of their mean values and take advantage of the additional constraining power from these moments (e.g., [88,89]). We leave this exploration to follow-up work.

-
- [1] D. N. Spergel, L. Verde, H. V. Peiris, E. Komatsu, M. R.olta, C. L. Bennett, M. Halpern, G. Hinshaw, N. Jarosik, A. Kogut *et al.*, *Astrophys. J. Suppl. Ser.* **148**, 175 (2003).
 - [2] N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo *et al.* (Planck Collaboration), *Astron. Astrophys.* **641**, A6 (2020).
 - [3] Y. B. Zeldovich and R. A. Sunyaev, *Astrophys. Space Sci.* **4**, 301 (1969).
 - [4] R. A. Sunyaev and Y. B. Zeldovich, *Astrophys. Space Sci.* **7**, 3 (1970).
 - [5] N. Kaiser, *Mon. Not. R. Astron. Soc.* **222**, 323 (1986).
 - [6] Z. Haiman, J. J. Mohr, and G. P. Holder, *Astrophys. J.* **553**, 545 (2001).
 - [7] S. W. Allen, A. E. Evrard, and A. B. Mantz, *Annu. Rev. Astron. Astrophys.* **49**, 409 (2011).
 - [8] M. Hasselfield, M. Hilton, T. A. Marriage, G. E. Addison, L. F. Barrientos, N. Battaglia, E. S. Battistelli, J. R. Bond, D. Crichton, S. Das *et al.*, *J. Cosmol. Astropart. Phys.* **07** (2013) 008.
 - [9] P. A. R. Ade, N. Aghanim, C. Armitage-Caplan, M. Arnaud, M. Ashdown, F. Atrio-Barandela, J. Aumont, C. Baccigalupi, A. J. Banday *et al.* Planck Collaboration, *Astron. Astrophys.* **571**, A20 (2014).
 - [10] Í. Zubeldia and A. Challinor, *Mon. Not. R. Astron. Soc.* **489**, 401 (2019).
 - [11] S. Bocquet, S. Grandis, L. E. Bleem, M. Klein, J. J. Mohr, T. Schrabback, T. M. C. Abbott, P. A. R. Ade, M. Aguena, A. Alarcon *et al.*, *Phys. Rev. D* **110**, 083510 (2024).
 - [12] E. Komatsu and U. Seljak, *Mon. Not. R. Astron. Soc.* **336**, 1256 (2002).
 - [13] S. Bhattacharya, D. Nagai, L. Shaw, T. Crawford, and G. P. Holder, *Astrophys. J.* **760**, 5 (2012).
 - [14] M. J. Wilson, B. D. Sherwin, J. C. Hill, G. Addison, N. Battaglia, J. R. Bond, S. Das, M. J. Devlin, J. Dunkley, R. Dünner *et al.*, *Phys. Rev. D* **86**, 122005 (2012).
 - [15] J. C. Hill and E. Pajer, *Phys. Rev. D* **88**, 063526 (2013).
 - [16] J. C. Hill and B. D. Sherwin, *Phys. Rev. D* **87**, 023527 (2013).
 - [17] P. A. R. Ade, N. Aghanim, C. Armitage-Caplan, M. Arnaud, M. Ashdown, F. Atrio-Barandela, J. Aumont, C. Baccigalupi, A. J. Banday *et al.* (Planck Collaboration), *Astron. Astrophys.* **571**, A21 (2014).
 - [18] J. C. Hill, B. D. Sherwin, K. M. Smith, G. E. Addison, N. Battaglia, E. S. Battistelli, J. R. Bond, E. Calabrese, M. J. Devlin, J. Dunkley *et al.*, *arXiv:1411.8004*.
 - [19] B. Horowitz and U. Seljak, *Mon. Not. R. Astron. Soc.* **469**, 394 (2017).
 - [20] G. Hurier and F. Lacasa, *Astron. Astrophys.* **604**, A71 (2017).
 - [21] L. Salvati, M. Douspis, and N. Aghanim, *Astron. Astrophys.* **614**, A13 (2018).
 - [22] B. Bolliet, B. Comis, E. Komatsu, and J. F. Macías-Pérez, *Mon. Not. R. Astron. Soc.* **477**, 4957 (2018).
 - [23] Planck Collaboration, P. A. R. Ade, N. Aghanim, C. Armitage-Caplan, M. Arnaud, M. Ashdown, F. Atrio-Barandela, J. Aumont, C. Baccigalupi, A. J. Banday *et al.*, *Astron. Astrophys.* **571**, A21 (2014).
 - [24] L. Thiele, J. C. Hill, and K. M. Smith, *Phys. Rev. D* **99**, 103511 (2019).
 - [25] D. Munshi, J. Smidt, S. Joudaki, and P. Coles, *Mon. Not. R. Astron. Soc.* **419**, 138 (2012).
 - [26] D. Munshi, P. Coles, and A. Heavens, *Mon. Not. R. Astron. Soc.* **428**, 2628 (2013).
 - [27] B. Bolliet, T. Brinckmann, J. Chluba, and J. Lesgourgues, *Mon. Not. R. Astron. Soc.* **497**, 1332 (2020).
 - [28] J. M. Kratochvil, Z. Haiman, and M. May, *Phys. Rev. D* **81**, 043519 (2010).
 - [29] J. M. Kratochvil, E. A. Lim, S. Wang, Z. Haiman, M. May, and K. Hufenberger, *Phys. Rev. D* **85**, 103513 (2012).
 - [30] A. Petri, Z. Haiman, L. Hui, M. May, and J. M. Kratochvil, *Phys. Rev. D* **88**, 123002 (2013).
 - [31] A. Petri, J. Liu, Z. Haiman, M. May, L. Hui, and J. M. Kratochvil, *Phys. Rev. D* **91**, 103511 (2015).

- [32] J. Liu, A. Petri, Z. Haiman, L. Hui, J. M. Kratochvil, and M. May, *Phys. Rev. D* **91**, 063507 (2015).
- [33] G. A. Marques, J. Liu, J. M. Zorrilla Matilla, Z. Haiman, A. Bernui, and C. P. Novaes, *J. Cosmol. Astropart. Phys.* **06** (2019) 019.
- [34] W. R. Coulton, J. Liu, I. G. McCarthy, and K. Osato, *Mon. Not. R. Astron. Soc.* **495**, 2531 (2020).
- [35] D. Zürcher, J. Fluri, R. Sgier, T. Kacprzak, and A. Refregier, *J. Cosmol. Astropart. Phys.* **01** (2021) 028.
- [36] M. Gatti, B. Jain, C. Chang, M. Raveri, D. Zürcher, L. Secco, L. Whiteway, N. Jeffrey, C. Doux, T. Kacprzak *et al.*, *Phys. Rev. D* **106**, 083509 (2022).
- [37] X. Liu, S. Yuan, C. Pan, T. Zhang, Q. Wang, and Z. Fan, *Mon. Not. R. Astron. Soc.* **519**, 594 (2022).
- [38] D. Lanzieri, F. Lanusse, C. Modi, B. Horowitz, J. Harnois-Déraps, and J.-L. Starck (LSST Dark EnergyScience (LSST DESC Collaboration), *Astron. Astrophys.* **679**, A61 (2023).
- [39] G. A. Marques, J. Liu, M. Shirasaki, L. Thiele, D. Grandón, K. M. Huffenberger, S. Cheng, J. Harnois-Déraps, K. Osato, and W. R. Coulton, *Mon. Not. R. Astron. Soc.* **528**, 4513 (2024).
- [40] C. P. Novaes, L. Thiele, J. Armijo, S. Cheng, J. A. Cowell, G. A. Marques, E. G. M. Ferreira, M. Shirasaki, K. Osato, and J. Liu, *Phys. Rev. D* **111**, 083510 (2025).
- [41] J. Armijo, G. A. Marques, C. P. Novaes, L. Thiele, J. A. Cowell, D. Grandón, M. Shirasaki, and J. Liu, *Mon. Not. R. Astron. Soc.* **537**, 3553 (2025).
- [42] C. Hikage, J. Schmalzing, T. Buchert, Y. Suto, I. Kayo, A. Taruya, M. S. Vogeley, F. Hoyle, I. Gott, J. Richard, and J. Brinkmann, *Publ. Astron. Soc. Jpn.* **55**, 911 (2003).
- [43] E. Komatsu, J. Dunkley, M. R. Nolte, C. L. Bennett, B. Gold, G. Hinshaw, N. Jarosik, D. Larson, M. Limon, L. Page *et al.*, *Astrophys. J. Suppl. Ser.* **180**, 330 (2009).
- [44] C. Hikage, T. Matsubara, P. Coles, M. Liguori, F. K. Hansen, and S. Matarrese, *Mon. Not. R. Astron. Soc.* **389**, 1439 (2008).
- [45] E. Komatsu and T. Kitayama, *Astrophys. J. Lett.* **526**, L1 (1999).
- [46] L. Thiele, J. C. Hill, and K. M. Smith, *Phys. Rev. D* **102**, 123545 (2020).
- [47] A. Sabyr, Z. Haiman, J. M. Z. Matilla, and T. Lu, *Phys. Rev. D* **105**, 023505 (2022).
- [48] J. L. Tinker, B. E. Robertson, A. V. Kravtsov, A. Klypin, M. S. Warren, G. Yepes, and S. Gottlöber, *Astrophys. J.* **724**, 878 (2010).
- [49] D. Blas, J. Lesgourgues, and T. Tram, *J. Cosmol. Astropart. Phys.* **07** (2011) 034.
- [50] N. Battaglia, J. R. Bond, C. Pfrommer, and J. L. Sievers, *Astrophys. J.* **758**, 75 (2012).
- [51] A. R. Duffy, J. Schaye, S. T. Kay, and C. Dalla Vecchia, *Mon. Not. R. Astron. Soc.* **390**, L64 (2008).
- [52] B. Bolliet, A. Kusiak, F. McCarthy, A. Sabyr, K. Surrao, J. C. Hill, J. Chluba, S. Ferraro, B. Hadzhiyska, D. Han *et al.*, *EPJ Web Conf.* **293**, 00008 (2024).
- [53] B. Bolliet, J. Colin Hill, S. Ferraro, A. Kusiak, and A. Krolewski, *J. Cosmol. Astropart. Phys.* **03** (2023) 039.
- [54] J. M. Zorrilla Matilla, Z. Haiman, D. Hsu, A. Gupta, and A. Petri, *Phys. Rev. D* **94**, 083506 (2016).
- [55] D. Coe, *arXiv:0906.4123*.
- [56] P. Ade, J. Aguirre, Z. Ahmed, S. Aiola, A. Ali, D. Alonso, M. A. Alvarez, K. Arnold, P. Ashton, J. Austermann *et al.*, *J. Cosmol. Astropart. Phys.* **02** (2019) 056.
- [57] K. Abazajian, G. Addison, P. Adshead, Z. Ahmed, S. W. Allen, D. Alonso, M. Alvarez, A. Anderson, K. S. Arnold, C. Baccigalupi *et al.*, *arXiv:1907.04473*.
- [58] A. Petri, *Astron. Comput.* **17**, 73 (2016).
- [59] H. Hadwiger, *Vorlesungen ueber Inhalt, Oberfläche und Isoperimetrie*, Die Grundlehren der mathematischen Wissenschaften (Springer, New York, 1957),
- [60] H. Minkowski, *Math. Ann.* **57**, 447 (1903).
- [61] J. Schmalzing, M. Kerscher, and T. Buchert, *Proc. Int. Sch. Phys. "Enrico Fermi"* **132**, 281 (1996).
- [62] F. Schaller, J. Wagner, and S. Kapfer, *J. Open Source Software* **5**, 2538 (2020).
- [63] T. Matsubara, *Phys. Rev. D* **81**, 083505 (2010).
- [64] H. Tomita, *Prog. Theor. Phys.* **76**, 952 (1986).
- [65] T. Matsubara, *arXiv:astro-ph/0006269*.
- [66] D. Munshi, L. van Waerbeke, J. Smidt, and P. Coles, *Mon. Not. R. Astron. Soc.* **419**, 536 (2012).
- [67] J. Carron, *Astron. Astrophys.* **551**, A88 (2013).
- [68] D. Kodwani, D. Alonso, and P. Ferreira, *Open J. Astrophys.* **2**, 3 (2019).
- [69] J. Hartlap, P. Simon, and P. Schneider, *Astron. Astrophys.* **464**, 399 (2007).
- [70] 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).
- [71] W. R. Coulton and B. D. Wandelt, *arXiv:2305.08994*.
- [72] W. R. Coulton, F. Villaescusa-Navarro, D. Jamieson, M. Baldi, G. Jung, D. Karagiannis, M. Liguori, L. Verde, and B. D. Wandelt, *Astrophys. J.* **943**, 178 (2023).
- [73] M. Hilton, C. Sifón, S. Naess, M. Madhavacheril, M. Oguri, E. Rozo, E. Rykoff, T. M. C. Abbott, S. Adhikari, M. Aguena *et al.*, *Astrophys. J. Suppl. Ser.* **253**, 3 (2021).
- [74] L. D. Shaw, O. Zahn, G. P. Holder, and O. Doré, *Astrophys. J.* **702**, 368 (2009).
- [75] K. Osato and M. Takada, *Phys. Rev. D* **103**, 063501 (2021).
- [76] A. Rotti, B. Bolliet, J. Chluba, and M. Remazeilles, *Mon. Not. R. Astron. Soc.* **503**, 5310 (2021).
- [77] S. K. Choi, J. Austermann, K. Basu, N. Battaglia, F. Bertoldi, D. T. Chung, N. F. Cothard, S. Duff, C. J. Duell, P. A. Gallardo *et al.*, *J. Low Temp. Phys.* **199**, 1089 (2020).
- [78] A. Gupta, J. M. Zorrilla Matilla, D. Hsu, and Z. Haiman, *Phys. Rev. D* **97**, 103515 (2018).
- [79] J. M. Zorrilla Matilla, Z. Haiman, D. Hsu, A. Gupta, and A. Petri, *Phys. Rev. D* **94**, 083506 (2016).
- [80] C.-A. Lin and M. Kilbinger, *Astron. Astrophys.* **583**, A70 (2015).
- [81] A. Diaz Rivero and C. Dvorkin, *Phys. Rev. D* **102**, 103507 (2020).
- [82] 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).

- [83] J. D. Hunter, *Comput. Sci. Eng.* **9**, 90 (2007).
- [84] C. C. Gramazio, D. H. Laidlaw, and K. B. Schloss, *IEEE Trans. Visual. Comput. Graph.* **23**, 521 (2017).
- [85] A. Sabyr, J. C. Hill, and Z. Haiman, tSZ Maps (2024), <https://columbialensing.github.io/#tSZ>.
- [86] A. Sabyr, tSZ_NG (2024), https://github.com/asabyr/tSZ_NG/.
- [87] D. Alonso, J. Sanchez, and A. Slosar (LSST Dark Energy Science Collaboration), *Mon. Not. R. Astron. Soc.* **484**, 4127 (2019).
- [88] M. C. Neyrinck, I. Szapudi, and A. S. Szalay, *Astrophys. J. Lett.* **698**, L90 (2009).
- [89] J. Carron and I. Szapudi, *Mon. Not. R. Astron. Soc.* **434**, 2961 (2013).