Reconstructing cosmic growth with kinetic Sunyaev-Zel'dovich observations in the era of stage IV experiments

David Alonso,[1] Thibaut Louis,[2] Philip Bull,[3,4] and Pedro G. Ferreira[1]

[1]*University of Oxford, Denys Wilkinson Building, Keble Road, Oxford, OX1 3RH, United Kingdom*
[2]*UPMC Univ Paris 06, UMR7095, Institut d'Astrophysique de Paris, F-75014 Paris, France*
[3]*California Institute of Technology, Pasadena, California 91125, USA*
[4]*Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, California, USA*

(Received 8 April 2016; published 22 August 2016)

Future ground-based cosmic microwave background (CMB) experiments will generate competitive large-scale structure data sets by precisely characterizing CMB secondary anisotropies over a large fraction of the sky. We describe a method for constraining the growth rate of structure to sub-1% precision out to $z \approx 1$, using a combination of galaxy cluster peculiar velocities measured using the kinetic Sunyaev-Zel'dovich (kSZ) effect, and the velocity field reconstructed from galaxy redshift surveys. We consider only thermal SZ-selected cluster samples, which will consist of $\mathcal{O}(10^4$–$10^5)$ sources for Stage 3 and 4 CMB experiments respectively. Three different methods for separating the kSZ effect from the primary CMB are compared, including a novel blind "constrained realization" method that improves signal-to-noise by a factor of ∼2 over a commonly-used aperture photometry technique. Assuming a correlation between the integrated tSZ y-parameter and the cluster optical depth, it should then be possible to break the kSZ velocity-optical depth degeneracy. The effects of including CMB polarization and SZ profile uncertainties are also considered. In the absence of systematics, a combination of future Stage 4 experiments should be able to measure the product of the growth and expansion rates, $\alpha \equiv fH$, to better than 1% in bins of $\Delta z = 0.1$ out to $z \approx 1$—competitive with contemporary redshift-space distortion constraints from galaxy surveys. We conclude with a discussion of the likely impact of various systematics.

DOI: 10.1103/PhysRevD.94.043522

I. INTRODUCTION

Galaxies and their big sisters, clusters, are test particles buffeted around by the cosmic gravitational field. If we could accurately measure their motions, as well as their positions, it would be possible to learn much more about the origin and evolution of large scale structure, the fundamental properties of gravity, and the constituents of the Universe. Measurements of large scale flows, or peculiar velocities, are complementary to other approaches to mapping out the Universe that use, for example, the cosmic microwave background, the distribution of galaxies, and weak gravitational lensing.

For the past few decades, there have been numerous attempts to embark on this somewhat quixotic enterprise. There are now peculiar velocity catalogues with between 10^3–10^4 objects, some of which span the whole celestial sphere, others that are deeper and more targeted [1–7]. It has been an arduous endeavour which, in some cases, has led to controversial results. Attempts at using direct distance indicators to galaxies or clusters (such as Tully-Fisher or fundamental plane relations) lead to shallow surveys with large uncertainties. Type Ia supernovae supply tighter constraints and allow for deeper surveys, but such surveys are, as yet, too sparse [8,9]; the same can be said of current kinetic Sunyaev-Zel'dovich measurements (the method that we will explore in this paper). On occasion, peculiar velocity surveys have led to results that are outliers within the standard cosmological canon: in the late 1980s they were used to argue for an $\Omega \sim 1$ universe [10], while in the 1990s and 2000s they were used to claim evidence for excessive bulk motion on large scales [11,12]. Given all this, and the rise of redshift space distortions (RSD) as a tool to learn about infall, direct measurements of peculiar velocities have become a neglected (and often maligned) area of research.

This is about to change. We are embarking on a new era of cosmological surveys in which we will map out the Universe with unprecedented precision. In particular, by mapping the cosmic microwave background (CMB) over vast swathes of sky with fine resolution and high sensitivity, it should be possible to construct a completely new class of peculiar velocity catalogues that may revolutionize the field. By measuring the scattering of CMB photons off moving free electrons, it is possible to pick up an effect—the kinetic Sunyaev-Zel'dovich (kSZ) effect—which is color blind (i.e. follows the CMB blackbody spectrum), and proportional to the bulk motion of the free electron density [13,14]. Understanding clusters as localized concentrations of free electrons, this effect can thus be used to make a *direct* measurement of the cluster peculiar velocity, independent of distance and redshift, which in the future could allow us to construct deep surveys of the large scale flows of the Universe.

The kSZ effect has already been detected statistically, arguably using the WMAP data [15], but most decisively with data from ACT, [16], Planck [17], ACTPol [18], and SPT [19], as well as the combined measurement of [20]. Pointed (i.e. single-cluster) detections also exist, e.g. [21]. The significance of the detections is still poor and not good enough to be able to extract cosmological information, but the outlook is promising. A number of experiments have ramped up their sensitivity and scope, most notably Advanced ACT and SPT-3G [22], and plans are under way to develop a consortium of telescopes, known as "Stage 4" (S4), that will allow us to construct definitive catalogues of kSZ peculiar velocity constraints with $\mathcal{O}(10^4$–$10^5)$ objects.

There have been a number of attempts at forecasting what might be possible with future kSZ catalogues [23–26]. Indeed, using such catalogues to constrain the pairwise streaming velocity or the velocity correlation tensor seems promising, leading to improvements by factors of up to a few in the dark energy figure of merit. These statistics probe larger scales, less contaminated by nonlinear growth and bias, and are complementary to more widely used clustering statistics in redshift space.

Even more promising is the idea of matching kSZ catalogues with density catalogues in such a way as to "divide out" the cosmic variance in the density/velocity field. The most likely velocity field can be reconstructed from a measurement of the 3D density field as traced by the number density of galaxies; one can then compare the reconstructed velocity field with the kSZ measurements and find constraints on a combination of the growth rate of structure and the cluster optical depth/ionization fraction. Adding in other measurements, it may even be possible to disentangle the two—making it possible to separately constrain cluster gas physics and the linear growth rate. The purpose of this paper is to explore this approach, unpacking the different steps that go into such an estimation, and assessing the various alternatives at each step. Crucial to our analysis is a realistic assessment of the uncertainties that should be ascribed to this method.

We structure the paper as follows. In Sec. II we describe the methods proposed to estimate the different ingredients of this procedure (the kSZ signal, the cluster optical depth, and the reconstructed velocities), as well as the forecasting formalism used. In Sec. III we compare three different kSZ measurement methods, and present the forecast constraints on the combination $\alpha \sim fH$ for each of them for several choices of current and next-generation CMB experiments and redshift surveys. Finally, in Sec. IV we summarize the results and discuss the advantages and limitations of the proposed approach.

II. GROWTH RECONSTRUCTION METHOD

The idea behind the method explored here is to match a reconstructed velocity field with CMB measurements of the kSZ effect to obtain a per-source measurement of the growth rate of structure. The potential of combining kSZ measurements with galaxy surveys has been discussed before: forecasts for combinations of upcoming experiments were explored in [23,27–29], and redshift surveys were essential in the first determination of the kSZ streaming velocity [16], as well as more recent attempts using the CMASS survey to pull out the kSZ signal at redshifts $z \sim 0.4$–0.7 [18]. In this section we build on previous work and lay out, in detail, the observables that we need to work with, and the various steps involved in building up a reliable estimator for the growth rate.

The fractional temperature fluctuations due to the thermal and kinetic Sunyaev-Zel'dovich effects are [14]

$$\left.\frac{\Delta \mathrm{T}}{\mathrm{T}}\right|_{\mathrm{tSZ}}(\nu,\hat{\mathbf{n}}) = f_{\mathrm{tSZ}}(\nu)\frac{\sigma_T}{m_e c^2}\int P_e(l_z,\hat{\mathbf{n}})dl_z \equiv f_{\mathrm{tSZ}}(\nu)y(\hat{\mathbf{n}}) \quad (1)$$

$$\left.\frac{\Delta \mathrm{T}}{\mathrm{T}}\right|_{\mathrm{kSZ}}(\hat{\mathbf{n}}) = -\sigma_T\int(\boldsymbol{\beta}\cdot\hat{\mathbf{n}})n_e(l_z,\hat{\mathbf{n}})dl_z \equiv -\beta_r\tau(\hat{\mathbf{n}}), \quad (2)$$

where n_e and $P_e = k_{\mathrm{B}}n_eT_e$ are the electron number density and pressure, σ_T is the Thomson scattering cross section, and $\beta_r \equiv \mathbf{v}\cdot\hat{\mathbf{n}}/c$ is the cluster's bulk velocity along the line of sight from the observer (parametrized by l_z). The spectral dependence of the thermal-Sunyaev-Zel'dovich (tSZ) effect is given by $f_{\mathrm{tSZ}}(\nu)$, and we have also defined the dimensionless Compton-y parameter, $y(\theta)$, and optical depth, $\tau(\theta)$, profiles as a function of angle from the center of the cluster (i.e. assuming sphericity). From Eq. (2), it is clear that a detection of the kSZ effect corresponds to a measurement of the combination $\beta_r \times \tau$; if an external estimate of τ can be made, this determines the local velocity field.

Let us now assume that we have a spectroscopic galaxy survey covering a volume that contains a number of tSZ-detected clusters, and that the redshifts of those clusters are known. As we will describe in Sec. II D, the galaxy distribution can be used to reconstruct the velocity field at the cluster positions up to a factor

$$\alpha(z) \equiv \frac{H(z)f(z)}{H_{\mathrm{fid}}(z)f_{\mathrm{fid}}(z)}, \quad (3)$$

where H and f are the expansion and growth rates, and the subscript "fid" labels quantities computed assuming the fiducial cosmology used to carry out the velocity reconstruction. For a given cosmology, the expected amplitude of the kSZ effect of a cluster i, as defined in Appendix, is $a^i_{\mathrm{kSZ}} = \beta^i_r\tau^i_{500}$. Assuming a value for τ_{500} and an estimate of the cluster's radial velocity, $\hat{\beta}_r$, from the velocity field reconstruction, we can sum over all clusters in a redshift interval $[z, z+\delta z]$ to obtain a likelihood for α,

$$-\log\mathcal{L} \equiv \chi^2(\alpha) = \sum_i \frac{(\alpha\hat{\beta}_r^i\tau_{500}^i - a_{\rm kSZ}^i)^2}{\mathcal{E}_i^2}. \quad (4)$$

Here, $a_{\rm kSZ}$ is the measured value of the kSZ amplitude, and $\mathcal{E}$ is the combined uncertainty in $\hat{\beta}_r$, τ_{500}, and $a_{\rm kSZ}$ for each cluster. Assuming that the errors on these parameters are independent and Gaussian-distributed, the uncertainty on α is given by

$$\sigma_\alpha^{-2} = \sum_i \mathcal{E}_i^{-2} \equiv \sum_i (\varepsilon^2_{a_{\rm kSZ},i} + \varepsilon^2_{\tau_{500},i} + \varepsilon^2_{\beta_r,i} + \varepsilon^2_{\tau_{500},i}\varepsilon^2_{\beta_r,i})^{-1}, \quad (5)$$

where $\varepsilon_x = \sigma_x/x$ are the relative uncertainties on the other three quantities.

The final uncertainty on α for a given combination of CMB experiment and spectroscopic survey depends on the number of clusters for which this process can be carried out. Both this, and the error on the measurement for each cluster, depend on the cluster halo mass, velocity, and redshift, and so we can rewrite Eq. (5) as an integral over their expected distributions,

$$\sigma_\alpha^{-2} = 4\pi f_{\rm sky}\frac{r^2(z)\delta z}{H(z)}\int_0^\infty dM\int_{-\infty}^\infty d\beta_r \times\frac{\tilde{\chi}(M,z)n(M,z)p(\beta_r|M,z)}{\mathcal{E}^2(M,\beta_r,z)}, \quad (6)$$

where $n(M,z)$ is the halo mass function (number density of dark matter haloes of mass $M\in[M,M+dM]$ in a given redshift interval), $p(\beta_r|M,z)$ is the distribution of halo radial velocities, and $\tilde{\chi}(M,z)$ is the detection efficiency for a cluster of a given mass and redshift for a given CMB experiment. The prefactor gives the volume of the redshift bin containing the clusters. For a given cosmology and set of survey specifications, we can therefore estimate the error on α by evaluating Eq. (6).

In what follows, we model the various measurement uncertainty terms in $\mathcal{E}(M,\beta_r,z)$ (Secs. II A and II D) and the detection efficiency $\tilde{\chi}(M,z)$ for tSZ-selected clusters (Sec. II C).

A. Cluster kSZ signal extraction

Unlike the tSZ effect, which has a distinctive spectral dependence, the kSZ effect has the same flat spectrum as the primary CMB—making the CMB anisotropies themselves an important source of contamination. Most kSZ detection methods therefore attempt to separate the two components by using differences in their angular distributions on the sky; while the angular extent of a typical galaxy cluster is of the order a few arcminutes (corresponding to $\ell\sim 3000$), the primary CMB anisotropies are strongly damped for $\ell\gtrsim 3000$, while dominating the power on much larger scales. An appropriately designed angular high-pass filter can therefore be used to separate the two contributions. This is the idea behind most kSZ extraction methods (e.g. see [30,31] and references therein). We compare three in this paper:

(i) The simplest is the *aperture photometry* (AP) filter, a blind method that uses a compensated circular filter with a radius similar to the cluster size to filter out the longer-wavelength CMB modes (Sec. II A 3).
(ii) An enhanced semiblind method, new to this work, that reconstructs and subtracts the CMB behind the aperture by using phase information from the surrounding area of sky, a technique known as *constrained realization* or in-painting (Sec. II A 2).
(iii) An optimal, minimum-variance *matched filter* estimator can be constructed by assuming a model for the spatial tSZ/kSZ profiles of the cluster. This entails making strong assumptions about the forms of $y(\theta)$ and $\tau(\theta)$, which leads to efficient filtering but potentially biased kSZ amplitude measurements (Sec. II A 1).

1. Matched filtering

Matched filtering [e.g. [32–34]] entails specifying a model for the spatial and spectral variation of the tSZ and kSZ signals, and then convolving the resulting set of filters with the (foreground-cleaned) maps. A perfectly matched filter will recover an unbiased estimate of the SZ amplitudes by strongly suppressing all other components with different spatial/spectral distributions. We model the data in each frequency band ν as

$$m_\nu(\hat{\mathbf{n}}) = \sum_i \mathrm{U}_\nu^i(\hat{\mathbf{n}})\cdot\mathbf{a}_i + n_\nu(\hat{\mathbf{n}}), \quad (7)$$

where $m_\nu(\hat{\mathbf{n}})\equiv\{\Delta T(\nu,\hat{\mathbf{n}})\}$ is the sky temperature measured in direction $\hat{\mathbf{n}}$, the noise term n_ν contains CMB, residual foreground (assumed zero here), and instrumental noise contributions, and the sum is over all clusters in the map. The matrix operator $\mathrm{U}_\nu(\hat{\mathbf{n}})\equiv(u_{\rm tSZ}(\nu,\hat{\mathbf{n}}),u_{\rm kSZ}(\nu,\hat{\mathbf{n}}))$ contains the tSZ and kSZ cluster spatial templates for each band, and $\mathbf{a}\equiv(a_{\rm tSZ},a_{\rm kSZ})$ is a vector of amplitude parameters (see Appendix for definitions and parametric profile models).

Assuming that the noise term is homogeneous, isotropic, and Gaussian-distributed [35], the log-likelihood for the amplitude parameters $\mathbf{a}$ is

$$\chi^2 = \int d^2l[\mathbf{m_l}-\mathbf{U_l}\cdot\mathbf{a}]^T\cdot\mathbf{C}_{\rm N}^{-1}(l)\cdot[\mathbf{m_l}-\mathbf{U_l}\cdot\mathbf{a}], \quad (8)$$

where $\mathbf{l}$ labels the flat-sky Fourier modes of the $\hat{\mathbf{n}}$-dependent quantities in the previous equations, and the various bold quantities are appropriately-constructed block

vectors and matrices containing the corresponding values for each Fourier mode/band/cluster. The total noise angular power spectrum is given by $\mathbf{C}_{\rm N}(l)$.

A minimum-variance estimate for $\mathbf{a}$ is then

$$\tilde{\mathbf{a}} \equiv \mathrm{Cov}(\tilde{\mathbf{a}}) \cdot \int d^2 l \mathbf{U}_{\mathbf{l}}^T \mathbf{C}_{\rm N}^{-1}(l) \mathbf{m}_{\mathbf{l}}, \tag{9}$$

with covariance

$$[\mathrm{Cov}(\tilde{\mathbf{a}})]^{-1} = \int d^2 l \mathbf{U}_{\mathbf{l}}^T \mathbf{C}_{\rm N}^{-1}(l) \mathbf{U}_{\mathbf{l}}. \tag{10}$$

As stated previously, we assume that the only relevant noise components are the primary CMB anisotropies and instrumental noise. The noise power spectrum is then

$$[\mathrm{C_N}(l)]_{\nu\nu'} = \mathrm{C}_l^{\rm CMB} + \frac{N_l^\nu}{(B_l^\nu)^2}\delta_{\nu\nu'}, \tag{11}$$

where N_l^ν and B_l^ν are the noise power spectrum and harmonic coefficients of the instrumental beam profile in frequency channel ν. In our fiducial analysis we will assume uncorrelated noise, so that $N_l^\nu = \sigma_{{\rm N},\nu}^2$, where $\sigma_{{\rm N},\nu}$ is the rms noise per steradian in each channel. Note that correlated instrumental noise (e.g. due to coherent atmospheric fluctuations) is expected to be non-negligible for actual ground-based experiments.

While matched filtering yields a minimum variance estimate of the kSZ amplitude, its effectiveness depends upon selecting the correct SZ profiles; otherwise, the estimates will be biased. Clusters are far from simple, ideal objects, however—profiles vary significantly between clusters, and the parametric profiles that are typically used tend to give only approximate fits to any given object. One could marginalize over the profile parameters, imposing a prior on them based on hydrodynamic simulations, for example, but even state of the art simulations fail to fit some features of real cluster samples. Matched filtering therefore necessitates a strong (and potentially unrealistic) prior to be placed on cluster physics, so substantial care must be exercised in the use of this technique.

2. Constrained realizations

While the exact shape of the mean kSZ cluster profile is currently very uncertain, we have precise information about the statistics of the primary CMB—its temperature power spectrum is modeled, and well-measured, out to high ℓ. This information can be used to construct and subtract a maximum-likelihood (ML) estimate of the CMB *behind* the cluster, without needing to assume a specific cluster model. The method for doing this, called constrained realization (CR) or "in-painting" of the CMB, has been used previously to fill-in masked regions of CMB maps, for example [e.g. [36–41]].

Begin by assuming that a cluster catalogue has been obtained, and a tSZ- and foreground-free map has been produced using a frequency-dependent filtering scheme. Our CR method then proceeds as follows:

(1) Define a disc $\mathcal{D}$ of radius θ_R around the center of each cluster that is large enough to encompass the bulk of the cluster's kSZ emission.
(2) Use the measured CMB fluctuations outside the disc to infer the ML value inside the disc (as described below).
(3) Subtract the maximum-likelihood estimate from inside the disc, and integrate the residual in the disc area to estimate the total kSZ flux.

The ML CMB temperature field, $\bar{T}_{\rm CMB}$, can be obtained by Wiener-filtering the (cleaned) map with the disc region $\mathcal{D}$ masked out [42]. The covariance of the residual CMB field, $T_{\rm CMB}^{\rm (true)} - \bar{T}_{\rm CMB}$, is given by $(\boldsymbol{C}^{-1} + \boldsymbol{N}^{-1})^{-1}$ [43], where $\boldsymbol{C}$ is the CMB covariance matrix (fixed to a best-fit power spectrum model), and $\boldsymbol{N}$ is the noise covariance matrix assuming infinite noise inside the disc,

$$\boldsymbol{N}_{ij}^{-1} = \begin{cases} \sigma_{\rm pix}^{-2}\delta_{ij} & \\ 0 & \text{if } i,j \in \mathcal{D} \end{cases}, \tag{12}$$

where $\sigma_{\rm pix}^2$ is the per-pixel noise variance (assumed homogeneous) of the tSZ-cleaned map outside the masked disc region. Our estimator for the kSZ flux is then

$$\hat{a}_{\rm kSZ}^{\rm CR} = \sum_{i\in\mathcal{D}} (m_i - \bar{T}_{i,\rm CMB})\Omega_{\rm pix} \tag{13}$$

where m_i is the value of the tSZ-filtered map in pixel i (with pixel area $\Omega_{\rm pix}$), and the sum is over all pixels inside $\mathcal{D}$. The variance of $\hat{a}_{\rm kSZ}^{\rm CR}$ is then given by

$$\mathrm{Var}(a_{\rm kSZ}^{\rm CR}) = \sum_{i,j\in\mathcal{D}} [(\boldsymbol{C}^{-1} + \boldsymbol{N}^{-1})_{ij}^{-1} + \sigma_{\rm pix}^2\delta_{ij}]\Omega_{\rm pix}^2. \tag{14}$$

The first term in square brackets is the variance for the reconstructed CMB from above, and the second is the instrumental noise variance (which would affect the kSZ term even in the absence of the CMB). We have assumed that the effect of the tSZ and foreground components is fully encapsulated in the enlarged noise variance of the tSZ-cleaned map.

The first term in Eq. (14) is difficult to evaluate, so we compute it as

$$\sum_{i,j\in\mathcal{D}} (\boldsymbol{C}^{-1} + \boldsymbol{N}^{-1})_{ij}^{-1} = \mathbf{u}^T \cdot (\boldsymbol{C}^{-1} + \boldsymbol{N}^{-1})^{-1} \cdot \mathbf{u}, \tag{15}$$

where $\mathbf{u}$ is a vector containing 1 in all pixels inside the disc, and 0 otherwise. The matrix inversion $(\boldsymbol{C}^{-1} + \boldsymbol{N}^{-1})^{-1} \cdot \mathbf{u}$ is carried out using a preconditioned conjugate gradient solver, with $\boldsymbol{C}^{-1}$ and $\boldsymbol{N}^{-1}$ applied in Fourier and real space

TABLE I. Specifications for representative CMB experiments. The reduced RMS noise levels used in the forecasts for the CR and AP methods for both experiments were 14 μK-arcmin (S3) and 1.75 μK-arcmin (S4).

Frequency (GHz)	Noise RMS (μK-arcmin)		Beam FWHM (arcmin)	
	S3	S4	S3	S4
28	78.0	9.8	7.1	14.0
41	71.0	8.9	4.8	10.0
90	7.8	1.0	2.2	5.0
150	6.9	0.9	1.3	2.8
230	25.0	3.1	0.9	2.0

respectively (where each matrix is sparsest) [44]. We checked that this estimate of the uncertainty was independent of pixel size.

This method has one free parameter: the choice of disc radius, θ_R. For this work, we chose θ_R to be such that 80% of the (expected) beam-convolved kSZ signal was enclosed in the disc. (The same criterion was used for the aperture photometry filter described in the next section.) The resulting flux estimate will therefore be biased, as some fraction of the signal will fall outside the disc; this bias must be corrected for analytically, or using simulations. A real analysis would presumably also compare several choices of θ_R to ensure stability of the results [cf. [18]].

Because it does not use information about the shape of the SZ profiles, the performance of this estimator is dictated by the noise level and the size of the disc that we consider. The CMB correlation function drops rapidly with separation angle, making the uncertainty on the residual $T_{\rm CMB} - \bar{T}_{\rm CMB}$ a steep function of the disc radius, θ_R. The uncertainty for clusters that subtend larger angles is therefore dominated by the CMB fluctuations, while the instrumental noise becomes more relevant for smaller discs.

We finish by noting that, both for this and the aperture photometry filter discussed in the next section, we assume that we have a single "reduced" tSZ-free CMB map, i.e. one in which the frequency channels have been combined together after filtering out the (frequency-dependent) tSZ signal. This is only possible for multiband experiments. The noise level of the reduced map was determined by assuming that all frequencies played some part in the removal of the tSZ signal, and so is obviously larger than the optimal noise level that would result from directly coadding all channels in the absence of tSZ. The reduced map noise levels used for different experiments are quoted in Table I.

3. Aperture photometry filter

Aperture photometry (AP) attempts to avoid specific assumptions about both the CMB statistics and cluster properties. It is conceptually similar to the constrained realization method from above, in that it tries to estimate and subtract the CMB fluctuations inside a disc centered around each cluster. Its "modeling" of the CMB is much simpler, however.

The method defines two concentric circles around each cluster, with radii θ_R and $\sqrt{2}\theta_R$ respectively, such that the areas of the inner and outer regions are the same. If CMB fluctuations have a typical angular size much larger than θ_R, they will be almost constant over the aperture. Subtracting the flux integrated over the outer region from the inner region will therefore result in zero mean CMB signal. Assuming that θ_R has been chosen such that most of the kSZ flux is inside the inner region, the integral of the residual there will be a good estimate of the total kSZ flux.

The simplicity of the AP method makes it possible to evaluate its performance analytically. The estimated kSZ contribution inside the inner region is

$$\tilde{\Delta}_{\rm kSZ}(\theta,\phi) = m(\theta,\phi) - \frac{1}{\pi\theta_R^2}\int_0^{2\pi} \times d\phi' \int_{\theta_R}^{\sqrt{2}\theta_R} d\theta'\theta' m(\theta',\phi'),$$

where $m(\hat{\mathbf{n}})$ is the tSZ-cleaned map, and (θ,ϕ) are cylindrical coordinates defined with respect to the center of the aperture. The total kSZ flux in the inner region is

$$a_{\rm kSZ}^{\rm AP} = \int_0^{2\pi} d\phi \int_0^{\theta_R} d\theta\theta \tilde{\Delta}_{\rm kSZ}(\theta,\phi) = \int_0^{2\pi} d\phi \int_0^{\infty} d\theta\theta W_{\rm AP}(\theta|\theta_R) m(\theta,\phi), \quad (16)$$

where the AP window function $W_{\rm AP}(\theta|\theta_R)$ is 1 for $0<\theta<\theta_R$, -1 for $\theta_R<\theta<\sqrt{2}\theta_R$, and 0 otherwise. For homogeneous and isotropic noise, the variance of $a_{\rm kSZ}^{\rm AP}$ is

$$\mathrm{Var}(a_{\rm kSZ}^{\rm AP}) = 2\pi\theta_R^4 \int_0^{\infty} dl l C_N(l) |\tilde{W}_{\rm AP}(l|\theta_R)|^2, \quad (17)$$

where $C_N(l)$ is the noise power spectrum (including CMB and instrumental noise), and $\tilde{W}_{\rm AP}$ is the Fourier transform of the AP filter, given by

$$\tilde{W}_{\rm AP}(l|\theta_R) = \frac{2J_1(l\theta_R) - \sqrt{2}J_1(\sqrt{2}l\theta_R)}{l\theta_R}. \quad (18)$$

We also validated this calculation numerically, using Gaussian realizations of the CMB.

While this method is in some sense model-independent, it is also biased (like the CR method, above), and has higher variance. The latter is a consequence of the nonvanishing primary CMB power on scales of order the aperture size; while suppressed due to Silk damping, CMB anisotropies

still dominate the kSZ signal on the typical (arcminute) angular scales of clusters. These contributions are not filtered by the AP method, and contribute significantly to the variance.

B. Cluster optical depth

As discussed above [see Eq. (2)], the kSZ flux measures a degenerate combination of optical depth and velocity, and so additional information is needed to recover the velocities by themselves. This can be achieved through a number of different methods—for example, the mean optical depth as a function of cluster mass and redshift can be calibrated using simulations [46], or from CMB polarization data [47]. One can also independently estimate τ by self-consistently modeling the ionized gas profile, or combining X-ray and tSZ information [48].

Following the results of [49] using hydrodynamical simulations, we have assumed a logarithmic scaling relation between the mean optical depth and the integrated Compton-y parameter:

$$\log_{10}\bar{\tau}_{500} = A + B\log_{10}\bar{Y}_{500}, \tag{19}$$

and that the value of τ_{500} for individual clusters will be scattered around this relation with a dispersion $\Delta\tau/\tau$. The relative uncertainty on τ_{500} is then

$$\varepsilon_\tau = \sqrt{B^2\left(\frac{\sigma_Y}{Y_{500}}\right)^2 + \left(\frac{\Delta\tau}{\tau}\right)^2}, \tag{20}$$

where σ_Y is the statistical uncertainty in the measurement of Y_{500}, given in Eq. (10), and we have assumed a scatter $\Delta\tau/\tau = 0.15$, in agreement with simulations [46]. Note that this value corresponds to the scatter in τ for a given mass range, and not the scatter around the $Y_{500} - \tau_{500}$ relation. In that sense Eq. (20) would conservatively overestimate the total uncertainty on τ_{500}.

Regarding the scaling parameter, B, here we have used the scaling of Y_{500} and τ_{500} with halo mass, M_{500}, according to the cluster models described in Appendix, to obtain

$$B = \frac{\alpha - 4/3}{\alpha - 2/3} \approx 0.41, \tag{21}$$

where $\alpha \simeq 1.79$ is the scaling exponent of the $Y_{500} - M_{500}$ relation [Eq. (24)]. This is in good agreement with the results of [49], who find a value of $B \simeq 0.48$.

Note that any systematic uncertainty in the $Y - \tau$ relation used in the analysis will not average down with the number of clusters in the sample, and would instead propagate directly into the final uncertainty on the growth rate (i.e. assuming that uncertainties in the relation are marginalized, to avoid biasing the measurement of α). At the moment, systematic uncertainties in the physical modeling of the intracluster medium are of the order of 10% [49], although there is hope that progress in our understanding of gas physics, as well as in the quality of future observational data, will reduce these uncertainties.

C. Detection efficiency of SZ-selected clusters

At this stage it is worth noting that, even though in principle this method could be used on cluster catalogs compiled using observables other than the thermal SZ flux, there are compelling reasons for using tSZ-selected clusters. As stated in the previous section, there are reasons to expect a correlation between the measured Compton-y parameter and the cluster optical depth. If that is the case, strong tSZ sources will on average also have large kSZ, and will therefore contribute significantly to the final growth constraints. We tested this assumption directly by running forecasts for SZ catalogs selected with signal-to-noise thresholds $q = 5$ and $q = 6$, finding essentially the same constraints on α. This shows that weak SZ sources have a negligible contribution in this method. Following this rationale, any cluster that does not yield a detectable tSZ signal (regardless of whether it has been selected optically or otherwise) would not contribute significantly to the final constraints. Furthermore, as described in the previous section, measurements of the tSZ flux could be crucial in constraining the cluster optical depth.

For a tSZ-selected cluster survey, the detection efficiency can be written as

$$\tilde{\chi}(M_{500}, z) = \int d(\ln Y_{500}^{\rm true}) \int_{q\sigma_N}^{\infty} dY_{500}^{\rm obs} \times P_{\rm SZ}(\ln Y_{500}^{\rm true}|M_{500}, z) P_{\rm det}(Y_{500}^{\rm obs}|Y_{500}^{\rm true}), \tag{22}$$

where $P_{\rm det}$ is the probability of obtaining a measurement $Y_{500}^{\rm obs}$ for a true integrated tSZ flux $Y_{500}^{\rm true}$ (see Appendix), and $P_{\rm SZ}$ is the distribution of integrated tSZ fluxes for clusters of mass M_{500} at redshift z, which accounts for the intrinsic scatter in the $Y - M$ relation. We have assumed a detection threshold of $q\sigma_N$, where σ_N is the noise on the measurement of Y_{500} (given by Eq. (10) for matched-filter detections), and q is the detection level above which clusters are accepted (e.g. $q = 5$ denotes a 5σ detection threshold).

Assuming Gaussian errors on the tSZ flux, the inner integral in Eq. (22) is

$$\int_{q\sigma_N}^{\infty} dY_{500}^{\rm obs} P_{\rm det}(Y_{500}^{\rm obs}|Y_{500}^{\rm true}) = \frac{1}{2}\left[1 + {\rm erf}\left[\frac{Y_{500}^{\rm true} - q\sigma_N}{\sqrt{2}\sigma_N}\right]\right].$$

The distribution of true tSZ fluxes is usually assumed to take a log-normal form,

$$P_{\rm SZ}(\ln Y_{500}|M_{500},z) = \frac{1}{\sqrt{2\pi}\sigma_{\ln Y_{500}}} \times \exp\left[-\frac{\ln^2(Y_{500}/\bar{Y}_{500})}{2\sigma^2_{\ln Y_{500}}}\right], \quad (23)$$

where $\bar{Y}_{500}(M_{500},z)$ and $\sigma_{\ln Y_{500}}$ are the mean and intrinsic scatter in the $Y-M$ relation. We adopt the empirical fitting function from [50], given by

$$\bar{Y}_{500} = Y_*\left[\frac{d_A(z)}{100\ {\rm Mpc}/h}\right]^{-2}\left[\frac{(1-b)M_{500}}{10^{14}M_\odot/h}\right]^{\alpha}E^{\beta}(z), \quad (24)$$

where d_A is the angular diameter distance, $1-b=0.8$, $Y_* = 2.42\times10^{-10}$ sr^2, $\alpha = 1.79\pm0.08$, $\beta = 0.66\pm0.5$, $E(z)\equiv H(z)/H_0$, and $\sigma_{\ln Y_{500}} = 0.127\pm0.023$.

D. Galaxy survey velocity reconstruction

In Newtonian theory, the relationship between the velocity and density fields is fully described by three nonlinear equations: the continuity, Euler, and Poisson equations [51]. The continuity equation reads

$$\dot{\delta} + \frac{1}{a}\nabla\cdot((1+\delta)\mathbf{v}) = 0. \quad (25)$$

While evaluating the time derivative $\dot{\delta}$ in general requires solving the nonlinear system of equations in full, the density field grows self-similarly ($\delta(t,\mathbf{x}) = D(t)\delta(t_0,\mathbf{x})$) in linear theory. After linearization, this allows us to rewrite the (Fourier space) continuity equation as

$$\mathbf{v}(t,\mathbf{k}) = \frac{Hf}{a}\frac{i\mathbf{k}}{k^2}\delta(t,\mathbf{k}), \quad (26)$$

where a is the scale factor, and $H\equiv\dot{a}/a$ and $f\equiv\dot{D}/D$ are the expansion and growth rates. A measurement of the three-dimensional density field can therefore be used to infer the velocity field on linear scales. In practice, this can be achieved by using the number counts from a spectroscopic galaxy survey as a (biased) proxy for the true density. Several sources of systematic uncertainties must be addressed, however:

(a) *Nonlinearities*: Eq. (26) is only valid in the linear regime; nonlinearities may introduce a bias in the recovered velocities. The impact of this effect can be mitigated by filtering out the smallest nonlinear scales, at the cost of introducing extra variance in the reconstructed velocity field. We provide a more quantitative description of these effects below.

(b) *Galaxy bias*: The relation between the observed galaxy number density and the true matter density field must be correctly modelled in order to avoid a biased reconstructed velocity field. While the connection between both fields has been shown to be well described by a linear, deterministic, and scale-independent bias factor, $\delta_{\rm gal} = b_g\delta$, on large scales, possible deviations from this model on small scales are a potentially dangerous systematic uncertainty.

(c) *Shot noise*: Noise due to a low number density of detected galaxies can significantly increase the variance of the reconstructed velocity field. A Wiener-filtering approach can be used to down-weight shot noise-dominated scales [52]. A key assumption of this method is the fact that all the galaxies in the spectroscopic survey would be used in the reconstruction of the velocity field, and not only those objects included in the tSZ catalog. This reduces the effect of shot noise and should bring the true reconstruction uncertainties closer to our estimate from N-body simulations outlined below.

(d) *Redshift-space distortions*: The nonzero radial peculiar velocities of galaxies modify their apparent redshift, and hence distort the recovered density field in an anisotropic manner. Redshift-space distortions are, however, well understood in linear theory, and can be fully incorporated into Eq. (26). An incorrect modeling of nonlinear RSDs could introduce important systematics in the reconstructed velocity field, however.

We obtained a best-case estimate of how accurately the true cluster velocities can realistically be recovered from the reconstructed velocity field by running a simple reconstruction algorithm on a suite of N-body simulations. The simulations were carried out using Gadget-2 [53], a tree-PM gravitational solver, which was run on initial conditions generated using second-order Lagrangian perturbation theory [54] at $z=49$ [55]. Each simulation contained 512^3 dark matter particles in a box of size $L_{\rm box} = 1400\ h^{-1}$ Mpc. A ΛCDM cosmological model was used, with parameters $(\Omega_M,\Omega_\Lambda,\Omega_b,h,\sigma_8,n_s) = (0.315, 0.685, 0.049, 0.67, 0.84, 0.96)$, compatible with the latest constraints from Planck [56]. Snapshots were output at redshifts $z=0$, 0.05, 0.15, 0.3 and 1, and dark matter haloes found in each of them using a friends-of-friends algorithm [57] with linking length $b_{\rm FOF}=0.2$.

For each snapshot, we estimated the reconstructed velocity for each halo as follows:

(1) The density field is estimated on a Cartesian grid of size $N_{\rm grid}=512$ using a Cloud-In-Cell algorithm.

(2) The density field is then smoothed using a Gaussian filter with standard deviation R_G as the characteristic scale. We studied the dependence of the reconstructed velocity field on the choice of smoothing scale by repeating this step for $R_G = \{0, 0.5, 1, 2, 4, 8\}\ h^{-1}$ Mpc.

(3) The velocity field is then estimated from the smoothed density field by solving the linearized continuity equation in Fourier space [Eq. (26)].

(4) A reconstructed velocity is assigned to each halo by interpolating the velocity field to the halo position, using a trilinear interpolation scheme.

We then compute the relative error between the reconstructed and true halo velocities for each halo, and study its statistics as a function of redshift and smoothing scale in different mass bins.

For halo masses in the range of interest, we find that it is always possible to find a smoothing scale that yields an unbiased estimate of the halo velocity, as well as roughly attaining minimum variance. Figure 1 shows this explicitly for the $z = 0.3$ snapshot. In all cases, we found the optimal smoothing scale to be in the range $R_G \in (2, 6)\ h^{-1}\,\mathrm{Mpc}$. The relative error in the reconstructed radial velocities is ∼50% across all masses and redshifts. We thus use the mean value $\varepsilon_{\beta_r} = 0.51$.

This estimate of the relative error due to the velocity reconstruction includes the contribution from nonlinear scales, but none of the other three effects listed above (galaxy bias, shot noise, and RSDs). A thorough evaluation of these lies beyond the scope of this paper. In any case, as evidenced by the results shown in Sec. III C, the uncertainty in the measured kSZ amplitude for each cluster should dominate the combined total uncertainty of the method [Eq. (5)], so we do not expect these caveats to significantly affect our results.

We also used the halo catalogues from these simulations to estimate the distribution of radial velocities $p(\beta_r|M, z)$ that enters Eq. (6). We find that $v_r \equiv c\beta_r$ is approximately Gaussian-distributed, with zero mean and a standard deviation given by

$$\sigma_v(M, z) \simeq \sigma_0(1+z)^{\gamma_0} - \sigma_1(1+z)^{\gamma_1}\log_{10} \times \left(\frac{M}{10^{14}\ h^{-1}\,M_\odot}\right),$$

with $(\sigma_0, \sigma_1) = (312 \pm 2, 22 \pm 3.5)$ km/s and $(\gamma_0, \gamma_1) = (0.87 \pm 0.01, 1.05 \pm 0.4)$. The uncertainties on the values of these parameters were estimated from the scatter across different simulations. It is worth noting that a mass and redshift dependence of the halo velocity dispersion is to be expected, and can be interpreted as follows: haloes form preferentially in higher density regions, and the typical range of allowed halo masses depends on that density (e.g. high-mass haloes are less likely to form in lower-density regions). The bulk velocity field is also directly correlated with the density field in the region; this is the basis of velocity reconstruction methods. Thus, assuming that haloes follow the overall matter velocity field, one expects the statistics of the halo velocities to also depend on halo mass [e.g. see [58]]. Note that this is different from a true halo bias in that, in this scenario, halo velocities do follow the total velocity field.

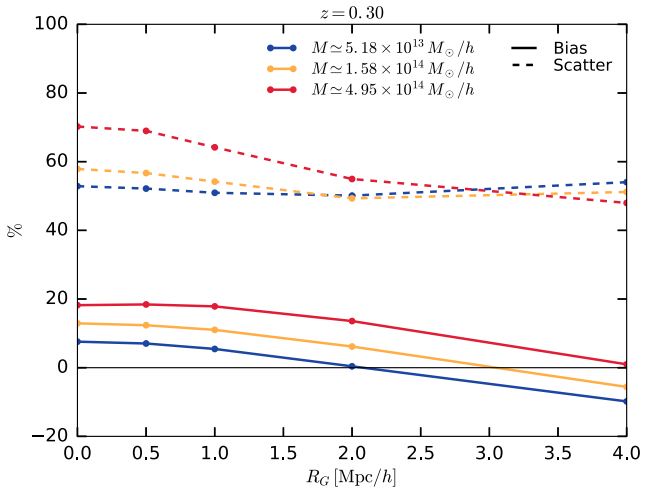

FIG. 1. Relative bias (solid lines) and standard deviation (dashed lines) of the reconstructed halo velocities for different Gaussian smoothing scales, in three mass bins at $z = 0.3$.

Finally, note that even though we have so far claimed that this method is able to yield a measurement of the quantity α defined in Eq. (3), since fH is the combination entering Eq. (26), the reconstructed velocity field is also sensitive to the normalization of the matter density field δ estimated from the galaxy overdensity. This relation is, on linear scales, determined by the galaxy bias b_g as well as the overall normalization of the density power spectrum which can be encoded in the parameter σ_8. Thus, in reality, this method measures the combination

$$\alpha \equiv \frac{f(z)H(z)b_g\sigma_8}{f_{\mathrm{fid}}H_{\mathrm{fid}}b_{g,\mathrm{fid}}\sigma_{8,\mathrm{fid}}}. \tag{27}$$

It should be possible to obtain tight priors on b_g and σ_8 from measurements of galaxy clustering and CMB power spectra, so we will regard α as mainly measuring the product fH in what follows. The existing uncertainties on these quantities should, however, be borne in mind.

III. RESULTS

In this section, we forecast how well each of the three kSZ extraction methods will be able to measure the expansion and growth rates using forthcoming Stage 3 and 4 ground-based CMB experiments.

A. Experimental setup

The current state of the art in CMB observation combines data sets from full-sky, space-based experiments (WMAP [59] and Planck [60]) with "Stage 2" ground-based experiments that focus on mapping the small-scale CMB anisotropies (e.g. ACTPol [61], SPT-Pol [62], and POLARBEAR [63]). Over the next few years, enhanced Stage 3 (S3) ground-based experiments (e.g. AdvACT [64] and SPT-3G [22]) will be rolled out, with larger numbers of detectors, multiple frequency channels, and the ability to survey a larger fraction of the sky. The high angular resolution and low noise levels of these experiments will make them ideal for cluster science, producing SZ

catalogues that contain $\mathcal{O}(10^4)$ sources over a wide range of masses and redshifts.

S3 experiments will eventually be superseded by a Stage 4 (S4) experiment, possibly composed of a set of ground-based facilities. Such an experiment would cover $\sim 20,000$ deg^2 on the sky, with noise levels of around 1 μK-arcmin. Such high sensitivity and large sky coverage is expected to increase the size of the corresponding cluster catalogue by at least an order of magnitude, making S4 an ideal experiment for the application of the method described here.

We consider a representative experimental specification for each Stage. For S3, we assume a wide ($f_{\rm sky} = 0.4$) survey with characteristics similar to those of AdvACT. The likely design of S4 is much less certain, so we consider an enhanced version of the S3 setup, with twice the beam width, eight times the sensitivity, and the same sky fraction. We assume Gaussian beams in every band for S3 and S4. The specifications for both experiments are detailed in Table I.

B. SZ catalogue properties

Using the formalism in Sec. II C, we predicted the expected mass and redshift distribution of the tSZ-selected cluster catalogues for each experiment (Fig. 2). To take foreground contamination into account, we eliminate the highest and lowest frequency channels for both experiments, assuming that they would be used as templates to remove synchrotron and dust contamination. Integrated mass and redshift distributions are shown in Fig. 3.

For both S3 and S4, we assumed a S/N threshold for cluster detection of $q = 6$ [65], yielding catalogues containing $\sim 10,000$ and $\sim 300,000$ sources respectively (in agreement with e.g. [64]). For S3, the bulk of the sample lies in the mass range $\log_{10} M_{500}/(h^{-1}M_\odot) \in (13.7, 14.5)$, and at redshifts $z \lesssim 0.6$, while S4 would be able to extend these ranges to $\log_{10} M_{500}/(h^{-1}M_\odot) \gtrsim 13.2$ and $z \lesssim 1.5$. We further validated this calculation by running our forecast pipeline with the specifications of the Planck survey [60], obtaining a catalogue with properties (mass and redshift distributions) similar to the one presented in [66]. In order to do this comparison, we used the same frequency channels included in [66], as well as the best-fit values of Ω_M and σ_8 found in their analysis. By varying the value of σ_8 within the 1σ uncertainty interval, we recovered cluster catalogs containing between ~ 220 and ~ 588 objects, in good agreement with the 439 sources found in [66].

Note that the average cluster size projected on the sky for S4 (given the size of the instrumental beam) is $\sim 4'$, while the expected number density of clusters for S4 is large (~ 35 deg^{-2}). It is straightforward to show that a fraction $f_{\rm blend} \approx 45\%$ of such a sample would overlap with other clusters on the sky ($f_{\rm blend} \approx 5\%$ for S3). Although the problem of cluster blending could in principle be overcome by using information about the cluster profiles, we have taken a conservative approach here and simply multiplied the number density of SZ sources by the expected fraction of nonoverlapping clusters, $1 - f_{\rm blend}$, essentially discarding the blended objects.

C. Comparison of kSZ extraction methods

We now compare the performance of the three different kSZ extraction methods described in Sec. II A: matched filtering (MF), constrained realizations (CR), and aperture photometry (AP).

Figure 4 shows the signal-to-noise ratio (SNR) of the kSZ amplitude measured for a set of characteristic cluster masses and redshifts, assuming a radial velocity $v_r = 300$ km/s. For a fixed mass, clusters subtend a larger angle on the sky with decreasing redshift, and the performance of the AP and CR methods is therefore significantly degraded at low z, as larger-scale CMB modes (which have larger variance) enter the filter region. This behavior is not

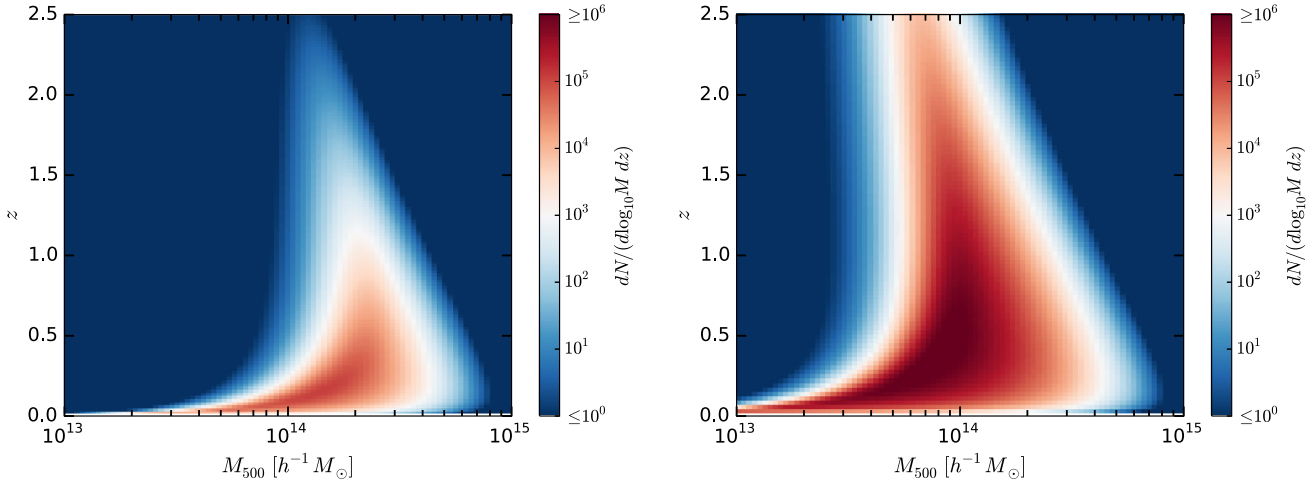

FIG. 2. Expected mass and redshift distributions for tSZ-selected clusters detected with the S3 (left) and S4 (right) experiments.

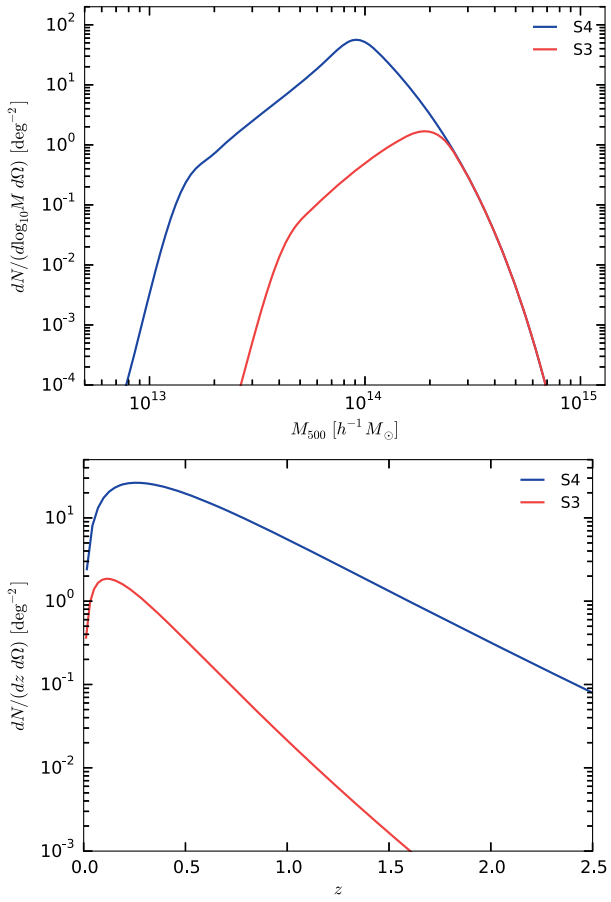

FIG. 3. Projected mass (top) and redshift (bottom) distributions of tSZ-selected clusters for S3 (red) and S4 (blue).

reproduced by the MF method, as knowledge of the SZ profile shape allows the cluster to be efficiently distinguished from CMB anisotropies, regardless of its increased variance. In fact, the MF method sees a slight increase in SNR at low redshift for low mass clusters, as the relatively weak signal can be added up coherently over a larger number of pixels.

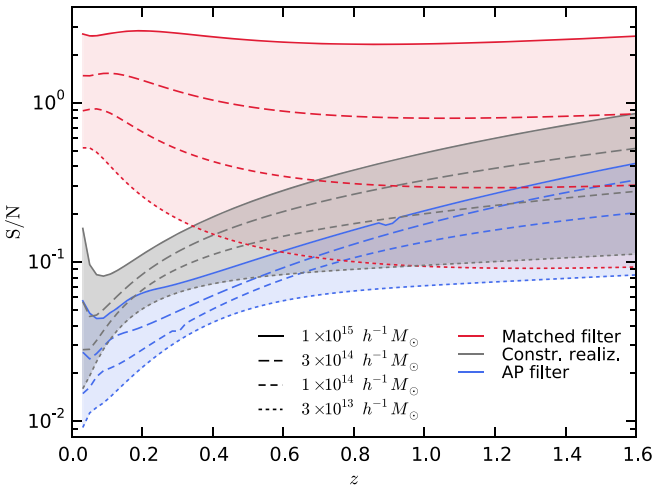

FIG. 4. Signal-to-noise ratios for the 3 different kSZ measurement methods, as a function of cluster redshift and mass.

The CR method shows a definite improvement over AP for all masses and redshifts, typically gaining a factor of ∼2 in SNR. While this is a factor of between 3–20 worse than the MF method, it is nevertheless a significant improvement for a model-independent method, especially considering the increased precision on the cosmological parameter measurement (see Sec. III D).

Profile uncertainty: While MF has by far the best performance in our simulations, its efficacy relies heavily on the accuracy of the assumed cluster profile. Given the current large uncertainty on the mean profile shape (e.g. see [18]), and the typical scatter in the profile from cluster to cluster, it is important to fold profile uncertainties into the errors on the recovered velocity. This is often achieved by repeating the analysis over a grid of profile parameter values for each cluster, although this rapidly becomes impractical as the number of parameters grows. Alternatively, a Monte Carlo parameter sampling approach can be taken, as discussed in [43].

The profile parameters are often poorly constrained however, and can suffer from strong degeneracies. We checked that this is likely to be the case by calculating Fisher matrices for the parameters of the GNFW profile (see Appendix for definitions), as constrained by the combined tSZ and kSZ profiles. There are several near-degeneracies in the Fisher matrix for S3 and S4, almost independent of redshift. After inversion, we find M_{500}, c_{500} (the concentration parameter), and γ (the outer slope) to be most strongly correlated with the cluster velocity, with correlation coefficients ranging from $|r| \simeq 0.6 - 0.9$ for a $10^{15}h^{-1}M_\odot$ cluster over a range of redshifts. Other parameter degeneracies make the matrix near-singular, however. Auxiliary information on the cluster shape (e.g. from X-ray observations or galaxy surveys) must therefore be added to break degeneracies in a real analysis. This typically relies on the use of scaling relations and simulations, the accuracy of which must also be folded into the uncertainty—in lieu of a generic procedure for doing this, we leave a quantitative analysis of profile shape uncertainties to future work.

Polarization: Both S3 and S4 are sensitive to polarization as well as total intensity. The tSZ and kSZ signals are expected to be almost completely unpolarized, while the CMB is not. Furthermore, the T and E CMB anisotropies are correlated, suggesting a possible way to improve the CMB reconstruction by including polarization information in the methods described in Sec. II A. As an example, we take the matched filter (MF) method and extend the profile matrix $\mathbf{U}$ in Eq. (10) with polarized channels in which the SZ profiles are set to zero. The variance of the kSZ amplitude is then computed as in Eq. (10), where the noise covariance now contains all auto- and cross-correlations between the temperature and polarization channels.

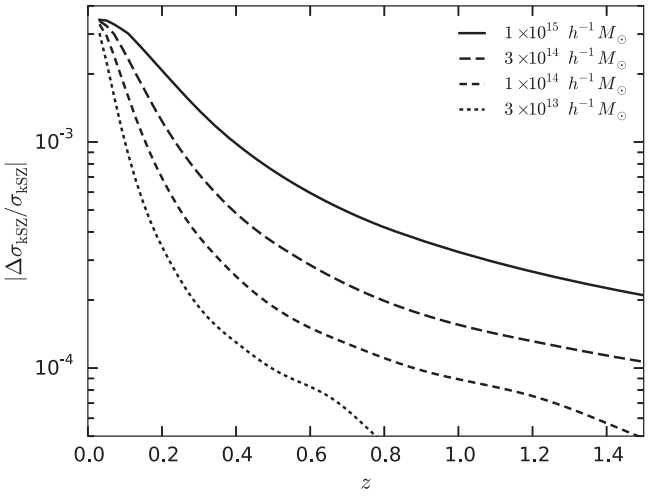

FIG. 5. Improvement in the kSZ measurement error for S3 after including polarization data. The marginal improvement for large clusters is due to the nonzero correlation between T and E, and is negligible overall.

Figure 5 shows the fractional change in the error on the kSZ measurement due to the inclusion of polarization information for S3. The improvement is negligible for all relevant cluster masses and redshifts. This result is disappointing but understandable: while the nonzero $T-E$ correlation does make it possible to better predict properties of the temperature field from the measured polarization field, the correlation is relatively small (∼10%), and basically negligible for noise-dominated scales (corresponding to most of the cluster sample).

S4 specification: Finally, as the specification of S4 is uncertain, it is worth exploring the possible benefit of different design strategies. For SZ cluster science, a narrower instrumental beam would allow the detection and characterization of less-massive and more-distant sources. Figure 6 shows the dependence of the kSZ uncertainty on the S4 beam FWHM for a $3\times10^{14}\ h^{-1}\,M_\odot$ cluster with $c\beta_r = 300$ km/s. Reducing the beam FWHM for S4 by a factor of ∼3 (i.e. from 3 arcmin to 1 arcmin) would improve the kSZ uncertainties by a similar factor at all redshifts for matched filters, and the uncertainty for the cluster-blind methods would gradually improve to a similar degree toward higher redshifts, where the smaller projected cluster size would benefit greatly from a reduced beam size. It is worth noting that, since the tSZ uncertainties would be similarly reduced, the effect on the performance of the method described here is twofold: first, it would increase the number of tSZ-detected clusters, and second, the kSZ uncertainties for those clusters would be reduced.

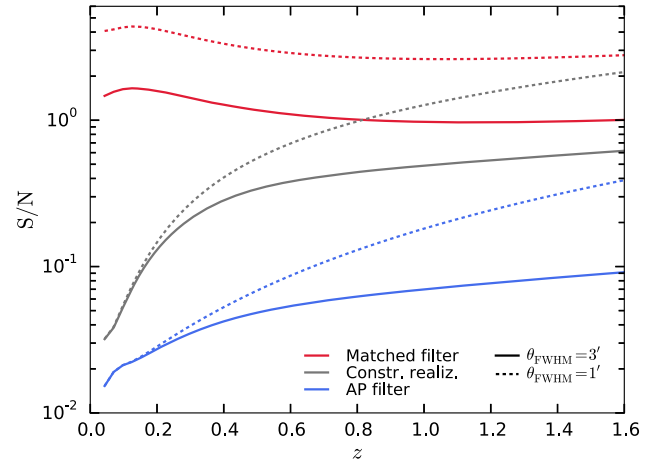

FIG. 6. kSZ signal-to-noise ratio for a cluster with mass $M_{500} = 3\times10^{14}\ h^{-1}\,M_\odot$ and radial velocity $v_r = 300$ km/s for S4 with two different beam FWHM: 3 arcminutes (solid lines) and 1 arcminute (dotted lines). The uncertainties for the matched filter approach improve by a factor ∼3 at all redshifts, while the improvement for constrained realizations and AP filtering improves gradually at higher redshifts, due to the smaller effective angle subtended by the cluster.

D. Expansion/growth rate constraints

We can now combine all of the information from the preceding sections to estimate the uncertainty on $\alpha \sim fH$ [Eq. (6)]. Figure 7 shows the forecast relative errors on α for the three kSZ extraction methods for both S3 and S4, using redshift bin widths $\Delta z = 0.2$ and 0.1 respectively, and assuming full overlap with a spectroscopic galaxy survey.

As expected, the matched filter (MF) method performs best, with S4 providing extremely competitive sub-1% measurements of fH out to $z = 1.5$. The blind CR method is only a factor of 2–3 worse above $z \approx 0.5$, which is also promising, while the AP method is a full order of magnitude down, mustering only ∼10% constraints for S4. The story for S3 is more one-sided, with the MF technique achieving ∼few% constraints out to $z \simeq 0.8$,

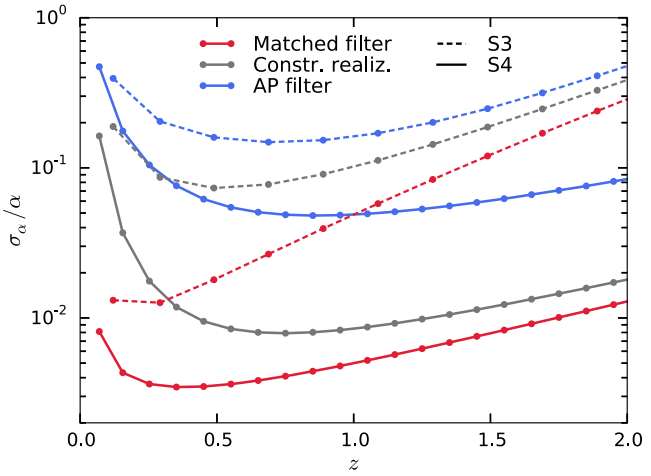

FIG. 7. Forecast constraints on $\alpha \sim fH$ for S3 (dashed lines) and S4 (solid), for the three different methods: matched filtering (MF; red), constrained realizations (CR; gray), and aperture photometry (AP; blue). The redshift bin widths are $\Delta z = 0.2$, 0.1 for S3 and S4 respectively, and we assume full redshift and area overlap with a spectroscopic survey.

while the CR and AP methods reach only ∼few × 10% at best. The MF method performs especially well at low redshift, where clusters can be well-resolved (especially by the high-resolution S3), allowing the shape information assumed by the filter to have the fullest effect. The difference is less pronounced at high z, so using blind methods here may be preferable due to their conservatism.

E. Dependence on galaxy survey overlap

The constraints on α ultimately depend on the availability of an overlapping spectroscopic galaxy redshift survey. To explore the importance of this issue, we selected three forthcoming galaxy surveys according to their expected time of completion: BOSS, DESI, and 4MOST. The forecast uncertainties for each of them are shown in Fig. 8, assuming the matched filter method for kSZ extraction. When estimating the overlap of these surveys with our model CMB experiments, we have assumed that both S3 and S4 will be southern hemisphere facilities. For comparison, the figure also includes the constraints for an "ideal" experiment, with full redshift and area overlap. Optimistic forecasts for fH from a Euclid-like spectroscopic galaxy survey, made by combining BAO and RSD Fisher forecasts from [67], are also shown for comparison.

The most competitive existing spectroscopic survey, in terms of surveyed volume, is SDSS-III's Baryon Oscillation Spectroscopic Survey (BOSS) [68]. The combination of its LOWZ and CMASS samples covers most of the redshift range out to $z = 0.7$ over ∼10,000 deg^2 on the sky, with a number density $n_g \sim 10^{-4}$ $(h^{-1}\,\mathrm{Mpc})^{-3}$. We assume a ∼50% area overlap (5,000 deg^2) between BOSS and our model S3 experiment, due to the northern hemisphere location of BOSS. The relatively low number density of sources in BOSS could severely affect the uncertainty on the reconstructed velocities, and so we conservatively doubled the size of the uncertainty ϵ_{β_r}. This is in agreement with the results of [18]. Using matched filters, an S3 experiment overlapping with BOSS could obtain a ∼7% measurement of fH in the range $0.1 < z < 0.7$, improving by a factor of ∼3 for S4.

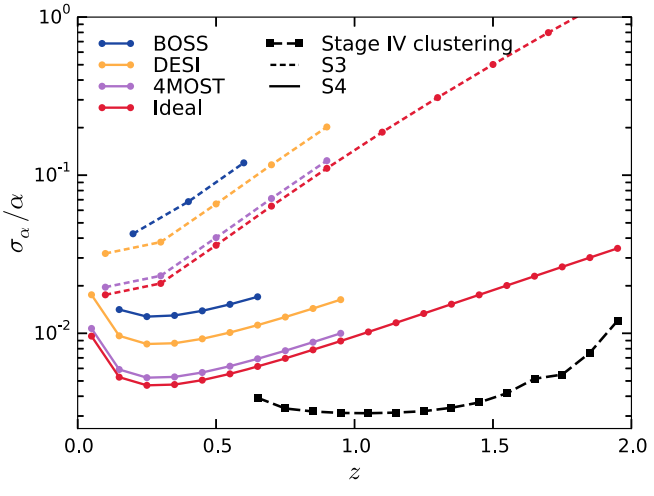

FIG. 8. Forecast constraints on $\alpha \sim fH$ for S3 and S4 (using the matched filter method), when three different galaxy surveys are used to provide the reconstructed velocity field. Results for an ideal (perfectly overlapping, sample variance-limited) survey are shown in red (cf. Fig. 7). Projected constraints on α from BAO + RSDs with a Euclid-like spectroscopic galaxy redshift survey are shown in black.

BOSS will be superseded by the Dark Energy Spectroscopic Instrument (DESI) [69], which will operate for 4 years starting in 2018. Jointly, its LRG and ELG samples will cover a similar fraction of the sky to BOSS, but now reaching out to $z \simeq 1.5$, and with a higher number density. We assume the same 50% overlap with the S3 and S4 surveys. Note that the number density will likely be too low to yield a reliable velocity field reconstruction in the high-z tail, and so we have only considered the redshift range $z < 1$ here. The larger number density and redshift coverage of DESI yields a small improvement in the forecast uncertainties on α compared to BOSS, with errors of ∼5–10% achievable with S3, improving by a factor of ∼4–10 for S4.

The 4-metre Multi-Object Spectroscopic Telescope (4MOST) [70] will carry out a similar spectroscopic survey to DESI in terms of area, depth, and number density, but from the southern hemisphere. Although 4MOST will not start operations until 2021, its near-total overlap with the survey areas of southern hemisphere CMB experiments such as AdvACT makes it ideal for this kind of analysis. We assumed an 80% area overlap (∼14,000 deg^2) with S3 and S4, and a redshift overlap for all $z < 1$. The main improvement over DESI lies in the larger area, which translates into a factor of ∼2 lower uncertainties on fH. This signal-to-noise level would make these measurements competitive with forecast RSD and BAO uncertainties for Stage IV galaxy surveys.

Finally we note that in the next decade, radio facilities such as the Square Kilometre Array (SKA) [71] will carry out spectroscopic galaxy surveys using the 21 cm radio line. Since any survey carried out by the SKA and its pathfinders would have almost complete overlap with both S3 and S4, it is worth exploring the constraints achievable by these surveys. Phase 1 of the SKA would be able to produce a 5,000 deg^2 survey with significant number densities out to $z \approx 0.4$ [72]. The constraints from this experiment would therefore be similar to those of DESI for this reduced redshift range. The survey would be extended during Phase 2 of SKA to cover ∼30,000 deg^2 out to $z \simeq 1.3$. The results for such a survey would be similar to those forecast for 4MOST.

IV. DISCUSSION

We have studied the potential of measuring the growth rate using a combination of a reconstructed velocity field

from a galaxy redshift survey and CMB observations of the kSZ effect. The performance of this approach depends on the uncertainties with which three quantities can be measured: the kSZ flux of each cluster, the cluster velocity reconstructed from the galaxy density field, and the cluster optical depth. Of these, we have found the kSZ measurement error to be the dominant source of statistical uncertainty for most redshifts and masses, and so we have delved deeper in the details of kSZ extraction.

To this end, we have discussed and compared three different methods to measure the kSZ with varying degrees of conservatism: matched filters (MF), which assume knowledge of both the CMB anisotropies and the mean cluster profiles; constrained realizations (CR), which only assume a model of the CMB statistics; and angular photometry filters (AP), which separate the primary CMB and kSZ components using only qualitative assumptions about their scale dependence.

We have shown that these assumptions have a critical effect on the resulting kSZ uncertainties: while AP filters cannot be used to obtain interesting constraints on $\alpha \sim fH$, constrained realizations could reduce the kSZ uncertainties significantly, yielding percent-level errors on this quantity assuming a perfectly overlapping galaxy redshift survey. Knowledge about the cluster profiles is necessary to reduce the uncertainties further, especially at low redshifts where clusters subtend larger solid angles. In this case, we have shown that with matched filtering, it would be possible to obtain kSZ errors small enough to make this method competitive with RSD-based measurements of the growth rate, which should yield subpercent uncertainties with Stage IV galaxy surveys. We have further shown how this method can be extended to make use of polarization data, although the level of improvement caused by the $T-E$ correlation in this case is negligible.

It is worth noting that the CR method we propose in this work, based on subtracting our best guess of the CMB anisotropies, can significantly improve the S/N of kSZ measurements compared with the commonly-used AP filter, but without requiring strong assumptions to be made about the shape of the cluster SZ profile (as is the case with MF). Although the CR method does require the CMB power spectrum to be specified, we are now at a point where it is known with sufficient precision to make this method practical.

The methods presented here build on a number of assumptions. While these should mostly be quite reasonable, it is worth bearing in mind the following caveats that will affect any future analysis with real data:

(i) The effectiveness of the matched filter technique depends strongly on the uncertainty in the assumed kSZ profile. Marginalizing over profile parameters (e.g. using MCMC sampling techniques) is difficult due to the strong degeneracies between most parameters, so high-quality external data (e.g. from X-ray and optical observations) is needed to better constrain the profile shapes.

(ii) The large number densities of clusters that will be detectable means that blending (overlapping clusters on the sky) will be an important problem—several tens of percent of clusters will be blended for S4. We have assumed that blended clusters can be identified and discarded.

(iii) We have ignored several potential biases and uncertainties in the velocity field reconstruction procedure, due to effects such as shot noise, RSDs, and non-linear and scale-dependent bias. Although the uncertainty on the kSZ measurements should dominate the overall error bar, the impact of these effects should be studied in depth. This is the subject of ongoing work.

(iv) We have ignored biases and contamination due to imperfect foreground subtraction. Some foregrounds (e.g. radio point sources, or the cosmic infrared background) are correlated with cluster positions, and so may not average down. It should, however, be possible to clean these foregrounds using their different frequency spectra.

(v) We have ignored sources of the kSZ effect that are not associated with clusters, such as the Ostriker-Vishniac effect from the diffuse IGM, and patchy kSZ from the epoch of reionization.

(vi) We have only quantified the statistical uncertainties in the three observables $(a_{\rm kSZ}, \beta_r, \tau_{500})$, neglecting any systematic errors in their measurement. The power of this method relies on averaging over many low-significance, single-cluster measurements of α by using large numbers of clusters. Systematic uncertainties do not average down however, and so, for a sufficiently large number of clusters, the method will eventually be dominated by them. This is particularly relevant for one of the key assumptions we have made: the existence of a well-calibrated $Y_{500}-\tau_{500}$ relationship, needed to break the $\tau-\beta_r$ degeneracy. Due to our imprecise current knowledge of cluster physics, systematic deviations can be expected at first [49], which will need to be correctly quantified.

An important aspect of this method is its different dependence on cosmic variance with respect to traditional clustering-based measurements of the growth rate. The statistics of a single realization of the density field can only be determined up to an accuracy defined by the number of modes accessible in a given survey region. This sample variance limit is easily reached by galaxy surveys, given a sufficiently high number density of sources. The performance of the method discussed here depends on different factors, however: the measurement errors ε_i, and the total number of SZ clusters for which this measurement can be carried out. The latter is, in turn, determined by the shape

and redshift dependence of the mass function and the total surveyed volume. Both sources of uncertainty can (in principle) be reduced without limit, by improving experimental parameters such as the noise sensitivity and angular resolution. This reduces the measurement uncertainties, and extends the mass range of the resulting cluster sample to smaller masses (although, for a fixed lower mass bound, the method will be limited by the number of haloes present in the surveyed patch, which is a different manifestation of the cosmic variance problem). This very fact also distinguishes this method from other procedures proposed in the literature to measure the kSZ effect, such as the pairwise kSZ signal [73] or the projected-field probe of [20]. These methods would in turn be less sensitive to cluster blending effects. It is also worth noting that the contribution to the total kSZ signal from lower-mass objects not included in the tSZ catalog, as well as the signal from the epoch of reionization would also have a sample-variance contribution to the final cluster kSZ signal. For the realistic noise levels explored in this paper, this contribution should be subdominant to that of the primary CMB, and we have neglected it.

This leads to an almost complete immunity to cosmic variance, which can be interpreted as follows: the parameter $\alpha \sim f(z)H(z)$ is measured from the combination of two different proxies for the *same* velocity field, and so the stochastic velocity terms essentially cancel out. This leaves behind a deterministic term that can be measured to arbitrary precision, limited only by the aforementioned sources of noise that go into the α estimator, and not by mode counting. A similar effect arises when two differently-biased tracers of the density field are combined to measure RSDs [74].

Due to their tSZ selection functions, and the choice of overlapping galaxy redshift surveys, the growth constraints from S3 and S4 will be mostly restricted to $z \lesssim 1$. This is exactly the regime in which the growth rate has the most to tell us, though—f deviates increasingly from unity at later times, when dark energy begins to dominate the expansion history. Precision measurements of both growth and expansion at these redshifts are vital to attempts to characterize dark energy and possible modifications of GR. The combination of the two, α, constrained by the method described here, is highly complementary to other combinations measured by probes such as BAO and RSDs. By probing the velocity field in a very different (and more direct) way, this method also provides a useful consistency check on RSDs, which use the 2D shape of the clustering pattern, and require a number of modeling assumptions. While a successful application of this method will need an excellent calibration of systematic uncertainties (especially those related to cluster gas physics), we have shown that combined kSZ and galaxy redshift survey analyses promise to become an important window into gravitational physics on large scales in the near future.

ACKNOWLEDGMENTS

We are grateful to Nicholas Battaglia, Jo Dunkley, Simone Ferraro, Sigurd Næss, and Emmanuel Schaan for useful comments and discussion. We also thank the anonymous referee, whose input improved the quality of the paper. D. A. is supported by the Beecroft Trust and ERC Grant No. 259505. T. L. is supported by ERC Grant No. 267117 (DARK) hosted by Universite Pierre et Marie Curie- Paris 6. P. B.'s research was supported by an appointment to the NASA Postdoctoral Program at the Jet Propulsion Laboratory, California Institute of Technology, administered by Universities Space Research Association under contract with NASA. P. G. F. acknowledges support from STFC, the Beecroft Trust and the Higgs Centre in Edinburgh.

APPENDIX: SZ PROFILES AND AMPLITUDES

In this appendix we describe the models that were used to estimate the amplitude and projected cluster profiles for the tSZ and kSZ components throughout this paper. Figure 9 shows examples of the profiles for two different cluster masses at $z = 0.3$.

1. Thermal SZ profile

The tSZ contribution to the CMB anisotropies is given by Eq. (1), with

$$f_{\rm tSZ}(\nu) \equiv \frac{q(e^q+1)}{e^q-1} - 4, \qquad q \equiv \frac{h\nu}{k_B T_{\rm CMB}}. \tag{A1}$$

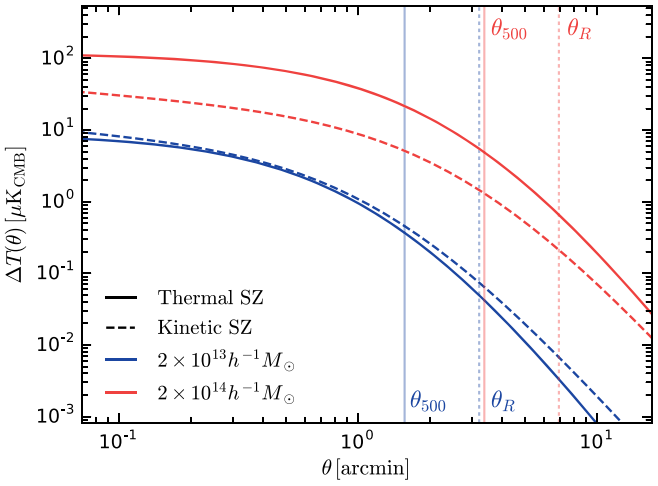

FIG. 9. The tSZ (solid lines) and kSZ (dashed) profiles for clusters with halo masses $2\times10^{13}\ h^{-1}\,M_\odot$ (blue) and $2\times10^{14}\ h^{-1}\,M_\odot$ (red), both at $z = 0.3$ with a radial velocity $v_r = -400$ km/s. The masses are chosen to be broadly representative of the S4 and S3 samples respectively (see Fig. 3). The vertical lines show the characteristic angular scale θ_{500} (solid), and the disc radius θ_R (dotted) for the AP filter for a Stage 3 experiment (1.4 arcmin beam), defined in Sec. II A 2.

To construct our model, we assume that the tSZ pressure profile is well described by the GNFW/Arnaud profile [75], i.e.

$$\frac{\sigma_T k_B}{m_e c^2} n_e(r) T_e(r) = L_0^{-1} p_p(r/R_{500}) \tag{A2}$$

$$p_p(x) = [(xc_{500})^\gamma [1 + (xc_{500})^\alpha]^{(\beta-\gamma)/\alpha}]^{-1}, \tag{A3}$$

where L_0 is a constant prefactor (with units of length), $p_p(x)$ is the dimensionless pressure profile, and the profile parameters are the best-fit values from [75]: $c_{500} = 1.156$, $\alpha = 1.062$, $\beta = 5.4807$, $\gamma = 0.3292$. Now, define the spherical aperture tSZ flux Y_{500} as

$$Y_{500} = \frac{4\pi}{d_A^2} \int_0^{R_{500}} dr r^2 n_e(r) \frac{k_B T_e}{m_e c^2} \sigma_T. \tag{A4}$$

Note that the spherical aperture flux is not a directly observable quantity, but can be related to the cylindrical aperture flux via the cluster model.

We can then write the tSZ anisotropy as in Eq. (7), with

$$a_{\rm tSZ} \equiv Y_{500}, \qquad u_{\rm tSZ}(\nu, \theta) = f_{\rm tSZ}(\nu) \frac{g_{\rm tSZ}(\theta/\theta_{500})}{4\pi\theta_{500}^2} \tag{A5}$$

$$g_{\rm tSZ}(x) \equiv \frac{\int_{-\infty}^{\infty} dx_z p_p(\sqrt{x_z^2 + x^2})}{\int_0^1 dx_r x_r^2 p_p(x_r)}, \tag{A6}$$

where x_r denotes the radius from the center of the cluster, and x_z is the distance along a line of sight through the cluster (with closest approach to the center, $x_z = 0$, at a radius x).

2. Kinetic SZ profile

The kSZ profile is determined by the electron density rather than the pressure profile. Here we will model n_e by assuming that the cluster gas is in hydrostatic equilibrium [76],

$$n_e(r) = \frac{\rho_{\rm gas}}{m_p \mu_e} = -\frac{r^2}{GM(<r) m_p \mu_e} \frac{dP_{\rm gas}}{dr}, \tag{A7}$$

where $\rho_{\rm gas}$ is the baryon mass density, $M(<r)$ is the total matter enclosed in a sphere of radius r, $P_{\rm gas}$ is the gas pressure (assumed thermal-only) and μ_e is the mean molecular weight per free electron. Assuming that the gas is fully ionized and has primordial composition, the thermal gas pressure is related to the electron pressure by $P_{\rm gas} = b_{\rm gas} P_e$, where $b_{\rm gas} \approx 1.93$ [77]. We assume a mass profile given by the Navarro-Frenk-White (NFW) universal halo profile [78],

$$M(<r) = M_{500} p_M(r/R_{500}), \tag{A8}$$

$$p(x) = \frac{\ln(1 + c_{500}x) - c_{500}x/(1 + c_{500}x)}{\ln(1 + c_{500}) - c_{500}/(1 + c_{500})}, \tag{A9}$$

and the Generalized-NFW (GNFW) pressure profile is given by $P_e(r) = m_e c^2/(L_0 \sigma_T) p_p(r/R_{500})$ as above. Evaluating Eq. (A7), we obtain

$$n_e(r) \equiv \frac{b_{\rm gas} m_e c^2}{G M_{500} m_p \mu_e \sigma_T} \frac{Y_{500} p_n(r/R_{500})}{4\pi\theta_{500}^2 \int_0^1 dx x^2 p_p(x)},$$

where we have defined the dimensionless number density profile $p_n(x) \equiv -x^2 p_p'(x)/p_M(x)$.

Now, define the spherical aperture optical depth τ_{500}, the quantity analogous to Y_{500}, as

$$\tau_{500} \equiv \frac{4\pi}{d_A^2(z)} \int_0^{R_{500}} dr r^2 n_e(r) \sigma_T \tag{A10}$$

$$= \frac{b_{\rm gas} m_e c^2 Y_{500} R_{500}}{G M_{500} m_p \mu_e} \frac{\int_0^1 dx x^2 p_n(x)}{\int_0^1 dx x^2 p_p(x)}$$

$$\approx 193 \frac{\int_0^1 dx x^2 p_n(x)}{\int_0^1 dx x^2 p_p(x)} \left[\frac{R_{500}}{1\ {\rm Mpc}/h}\right] \left[\frac{Y_{500}}{{\rm srad}^2}\right] \left[\frac{M_{500}}{10^{14}\ M_\odot/h}\right]^{-1}. \tag{A11}$$

The kSZ anisotropy can finally be written as in Eq. (7),

$$a_{\rm kSZ} \equiv -\beta_r \tau_{500}, \qquad u_{\rm kSZ}(\nu, \theta) = \frac{g_{\rm kSZ}(\theta/\theta_{500})}{4\pi\theta_{500}^2}, \tag{A12}$$

$$g_{\rm kSZ}(x) \equiv \frac{\int_{-\infty}^{\infty} dx_z p_n(\sqrt{x_z^2 + x^2})}{\int_0^1 dx_r x_r^2 p_n(x_r)}. \tag{A13}$$

Note that throughout this Appendix, we have described a single cluster model. The abundance of overlapping clusters along the line of sight was quantified in Sec. III B.

[1] J. L. Tonry, A. Dressler, J. P. Blakeslee, E. A. Ajhar, A. B. Fletcher, G. A. Luppino, M. R. Metzger, and C. B. Moore, Astrophys. J. **546**, 681 (2001).
[2] G. Wegner, M. Bernardi, C. N. A. Willmer, L. N. da Costa, M. V. Alonso, P. S. Pellegrini, M. A. G. Maia, O. L. Chaves, and C. Rité, Astron. J. **126**, 2268 (2003).
[3] C. M. Springob, K. L. Masters, M. P. Haynes, R. Giovanelli, and C. Marinoni, Astrophys. J. Suppl. Ser. **172**, 599 (2007).
[4] D. A. Dale, R. Giovanelli, M. P. Haynes, L. E. Campusano, E. Hardy, and S. Borgani, Astrophys. J. **510**, L11 (1999).
[5] M. Colless, R. P. Saglia, D. Burstein, R. L. Davies, R. K. McMahan, and G. Wegner, Mon. Not. R. Astron. Soc. **321**, 277 (2001).
[6] J. A. Willick, Astrophys. J. **522**, 647 (1999).
[7] A. Johnson, C. Blake, J. Koda, Y.-Z. Ma, M. Colless, M. Crocce, T. M. Davis, H. Jones, C. Magoulas, J. R. Lucey *et al.*, Mon. Not. R. Astron. Soc. **444**, 3926 (2014).
[8] A. G. Riess, W. H. Press, and R. P. Kirshner, Astrophys. J. **445**, L91 (1995).
[9] S. J. Turnbull, M. J. Hudson, H. A. Feldman, M. Hicken, R. P. Kirshner, and R. Watkins, Mon. Not. R. Astron. Soc. **420**, 447 (2012).
[10] A. Dekel, E. Bertschinger, A. Yahil, M. A. Strauss, M. Davis, and J. P. Huchra, Astrophys. J. **412**, 1 (1993).
[11] T. R. Lauer and M. Postman, Astrophys. J. **425**, 418 (1994).
[12] R. Watkins, H. A. Feldman, and M. J. Hudson, Mon. Not. R. Astron. Soc. **392**, 743 (2009).
[13] R. A. Sunyaev and Y. B. Zeldovich, Comments Astrophys. Space Phys. **4**, 173 (1972).
[14] R. A. Sunyaev and I. B. Zeldovich, Annu. Rev. Astron. Astrophys. **18**, 537 (1980).
[15] A. Kashlinsky, F. Atrio-Barandela, D. Kocevski, and H. Ebeling, Astrophys. J. **686**, L49 (2008).
[16] N. Hand *et al.*, Phys. Rev. Lett. **109**, 041101 (2012).
[17] P. A. R. Ade *et al.* (Planck Collaboration XXXVII), Astron. Astrophys. **586**, A140 (2016).
[18] E. Schaan *et al.*, Phys. Rev. D **93**, 082002 (2016).
[19] B. Soergel, S. Flender, K. T. Story, L. Bleem, T. Giannantonio, G. Efstathiou, E. Rykoff, B. A. Benson, T. Crawford, S. Dodelson *et al.*, Mon. Not. R. Astron. Soc. **461**, 3172 (2016).
[20] J. C. Hill, S. Ferraro, N. Battaglia, J. Liu, and D. N. Spergel, Phys. Rev. Lett. **117**, 051301 (2016).
[21] J. Sayers, T. Mroczkowski, M. Zemcov, P. M. Korngut, J. Bock, E. Bulbul, N. G. Czakon, E. Egami, S. R. Golwala, P. M. Koch *et al.*, Astrophys. J. **778**, 52 (2013).
[22] B. A. Benson, P. A. R. Ade, Z. Ahmed, S. W. Allen, K. Arnold, J. E. Austermann, A. N. Bender, L. E. Bleem, J. E. Carlstrom, C. L. Chang *et al.*, *Proceedings SPIE Int. Soc. Opt. Eng.* 9153, 91531P (2014).
[23] C. Hernandez-Monteagudo, L. Verde, R. Jimenez, and D. N. Spergel, Astrophys. J. **643**, 598 (2006).
[24] S. Bhattacharya and A. Kosowsky, Astrophys. J. **659**, L83 (2007).
[25] S. Bhattacharya and A. Kosowsky, Phys. Rev. D **77**, 083004 (2008).
[26] E.-M. Mueller, F. de Bernardis, R. Bean, and M. D. Niemack, Astrophys. J. **808**, 47 (2015).
[27] O. Dore, J. F. Hennawi, and D. N. Spergel, Astrophys. J. **606**, 46 (2004).
[28] S. DeDeo, D. N. Spergel, and H. Trac, astro-ph/0511060.
[29] J. C. Hill, S. Ferraro, N. Battaglia, J. Liu, and D. N. Spergel, Phys. Rev. Lett. **117**, 051301 (2016).
[30] O. Forni and N. Aghanim, Astron. Astrophys. **420**, 49 (2004).
[31] P. A. R. Ade *et al.* (Planck Collaboration XIII), Astron. Astrophys. **561**, A97 (2014).
[32] D. Herranz, J. L. Sanz, M. P. Hobson, R. B. Barreiro, J. M. Diego, E. Martínez-González, and A. N. Lasenby, Mon. Not. R. Astron. Soc. **336**, 1057 (2002).
[33] J.-B. Melin, J. G. Bartlett, and J. Delabrouille, Astron. Astrophys. **459**, 341 (2006).
[34] M. Li, R. E. Angulo, S. D. M. White, and J. Jasche, Mon. Not. R. Astron. Soc. **443**, 2311 (2014).
[35] Note that these assumptions can easily be relaxed in a practical situation.
[36] P. Abrial, Y. Moudden, J.-L. Starck, J. Fadili, J. Delabrouille, and M. K. Nguyen, Statistical Methodology **5**, 289 (2008).
[37] K. T. Inoue, P. Cabella, and E. Komatsu, Phys. Rev. D **77**, 123539 (2008).
[38] M. Bucher and T. Louis, Mon. Not. R. Astron. Soc. **424**, 1694 (2012).
[39] J. Kim, P. Naselsky, and N. Mandolesi, Astrophys. J. **750**, L9 (2012).
[40] F. Elsner and B. D. Wandelt, Astron. Astrophys. **549**, A111 (2013).
[41] D. S. Seljebotn, K.-A. Mardal, J. B. Jewell, H. K. Eriksen, and P. Bull, Astrophys. J. Suppl. Ser. **210**, 24 (2014).
[42] The Wiener-filter solution covers the full map, including the masked region.
[43] P. Bull, I. K. Wehus, H. K. Eriksen, P. G. Ferreira, U. Fuskeland, K. M. Górski, and J. B. Jewell, Astrophys. J. Suppl. Ser. **219**, 10 (2015).
[44] Note that several other fast methods for solving the Wiener filtering equation have been developed in the literature, e.g. [38,40,41,45].
[45] H. K. Eriksen, I. J. O'Dwyer, J. B. Jewell, B. D. Wandelt, D. L. Larson, K. M. Górski, S. Levin, A. J. Banday, and P. B. Lilje, Astrophys. J. Suppl. Ser. **155**, 227 (2004).
[46] N. Battaglia, J. R. Bond, C. Pfrommer, and J. L. Sievers, Astrophys. J. **806**, 43 (2015).
[47] S. Y. Sazonov and R. A. Sunyaev, Mon. Not. R. Astron. Soc. **310**, 765 (1999).
[48] N. Sehgal, A. Kosowsky, and G. Holder, Astrophys. J. **635**, 22 (2005).
[49] N. Battaglia, arXiv:1607.02442.
[50] P. A. R. Ade *et al.* (Planck Collaboration XXIV), arXiv:1502.01597.
[51] J. Binney and S. Tremaine, *Galactic Dynamics: Second Edition* (Princeton University Press, Princeton, NJ, 2008).
[52] J. D. Cohn, M. White, T.-C. Chang, G. Holder, N. Padmanabhan, and O. Doré, Mon. Not. R. Astron. Soc. **457**, 2068 (2016).
[53] V. Springel, Mon. Not. R. Astron. Soc. **364**, 1105 (2005).
[54] M. Crocce, S. Pueblas, and R. Scoccimarro, Mon. Not. R. Astron. Soc. **373**, 369 (2006).
[55] Code available at https://github.com/damonge/IC_DAM.
[56] P. A. R. Ade *et al.* (Planck Collaboration XIII), arXiv:1502.01589.

[57] Code available at https://github.com/damonge/MatchMaker.
[58] W. A. Hellwing, A. Barreira, C. S. Frenk, B. Li, and S. Cole, Phys. Rev. Lett. **112**, 221102 (2014).
[59] C. L. Bennett, D. Larson, J. L. Weiland, N. Jarosik, G. Hinshaw, N. Odegard, K. M. Smith, R. S. Hill, B. Gold, M. Halpern *et al.*, Astrophys. J. Suppl. Ser. **208**, 20 (2013).
[60] Planck Collaboration I, Astron. Astrophys. **571**, A1 (2014).
[61] S. Naess, M. Hasselfield, J. McMahon, M. D. Niemack, G. E. Addison, P. A. R. Ade, R. Allison, M. Amiri, N. Battaglia, J. A. Beall *et al.*, J. Cosmol. Astropart. Phys. 10 (2014) 007.
[62] R. Keisler, S. Hoover, N. Harrington, J. W. Henning, P. A. R. Ade, K. A. Aird, J. E. Austermann, J. A. Beall, A. N. Bender, B. A. Benson *et al.*, Astrophys. J. **807**, 151 (2015).
[63] Polarbear Collaboration, Astrophys. J. **794**, 171 (2014).
[64] S. W. Henderson, R. Allison, J. Austermann, T. Baildon, N. Battaglia, J. A. Beall, D. Becker, F. De Bernardis, J. R. Bond, E. Calabrese *et al.*, J. Low Temp. Phys. **184**, 772 (2016).
[65] Using a smaller signal-to-noise threshold would yield a larger catalogue, containing fainter and less massive objects. However, the kSZ signal-to-noise of those new objects is equally reduced, and their contribution to the final constraint on α is negligible. We verified this by repeating the analysis for $q = 4$.
[66] P. A. R. Ade *et al.* (Planck Collaboration XXXII), Astron. Astrophys. **581**, A14 (2015).
[67] P. Bull, Astrophys. J. **817**, 26 (2016).
[68] S. Alam, F. D. Albareti, C. Allende Prieto, F. Anders, S. F. Anderson, T. Anderton, B. H. Andrews, E. Armengaud, É. Aubourg, S. Bailey *et al.*, Astrophys. J. Suppl. Ser. **219**, 12 (2015).
[69] M. Levi, C. Bebek, T. Beers, R. Blum, R. Cahn, D. Eisenstein, B. Flaugher, K. Honscheid, R. Kron, O. Lahav *et al.*, arXiv:1308.0847.
[70] R. S. de Jong, S. Barden, O. Bellido-Tirado, J. Brynnel, C. Chiappini, É. Depagne, R. Haynes, D. Johl, D. P. Phillips, O. Schnurr *et al.*, Proc. SPIE Int. Soc. Opt. Eng. **9147**, 91470M (2014).
[71] R. Maartens, F. B. Abdalla, M. Jarvis, and M. G. Santos, *Proc Sci.*, AASKA14 (2015) 016 [arXiv:1501.04076].
[72] S. Yahya, P. Bull, M. G. Santos, M. Silva, R. Maartens, P. Okouma, and B. Bassett, Mon. Not. R. Astron. Soc. **450**, 2251 (2015).
[73] P. G. Ferreira, R. Juszkiewicz, H. A. Feldman, M. Davis, and A. H. Jaffe, Astrophys. J. **515**, L1 (1999).
[74] P. McDonald and U. Seljak, J. Cosmol. Astropart. Phys. 10 (2009) 007.
[75] M. Arnaud, G. W. Pratt, R. Piffaretti, H. Böhringer, J. H. Croston, and E. Pointecouteau, Astron. Astrophys. **517**, A92 (2010).
[76] X. Shi, E. Komatsu, D. Nagai, and E. T. Lau, Mon. Not. R. Astron. Soc. **455**, 2936 (2016).
[77] N. Battaglia, J. R. Bond, C. Pfrommer, and J. L. Sievers, Astrophys. J. **758**, 74 (2012).
[78] J. F. Navarro, C. S. Frenk, and S. D. M. White, Astrophys. J. **462**, 563 (1996).